

**NBA 6120**

---

Chip Making  
&  
Microprocessor Technology

Donald P. Greenberg

Lecture #2

August 31, 2015

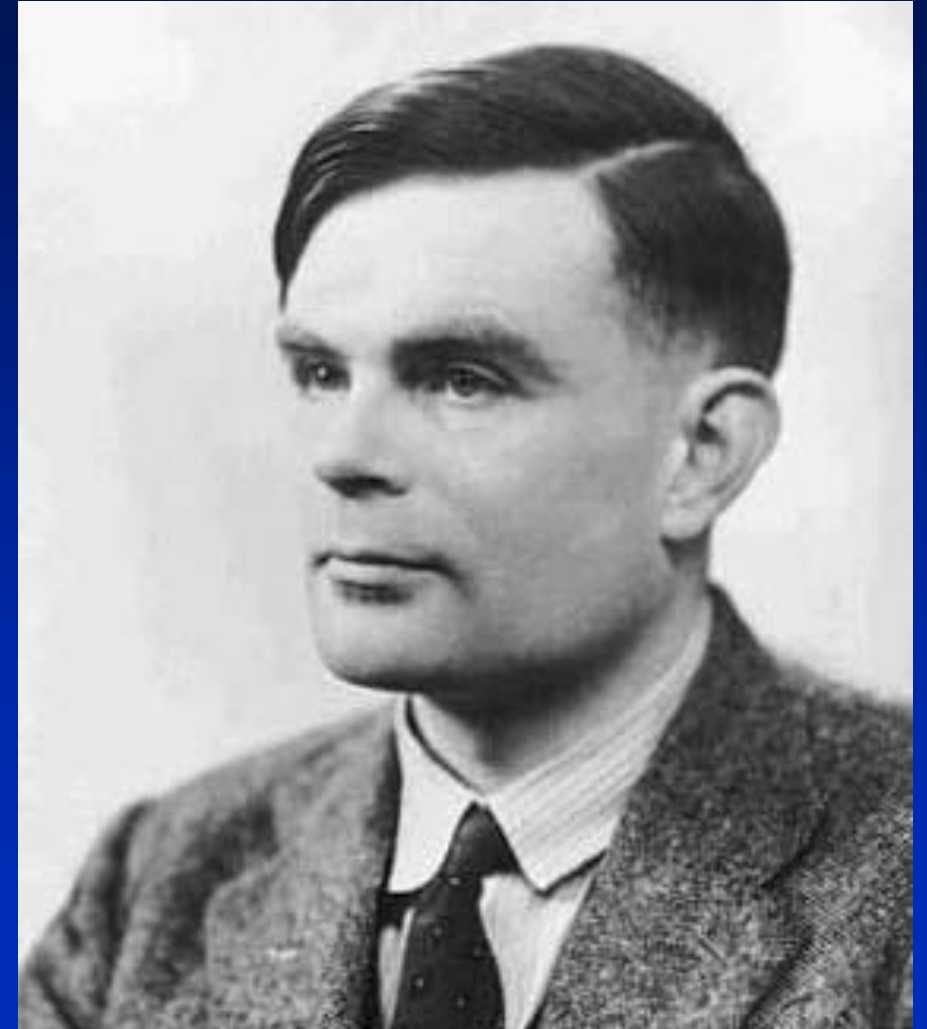
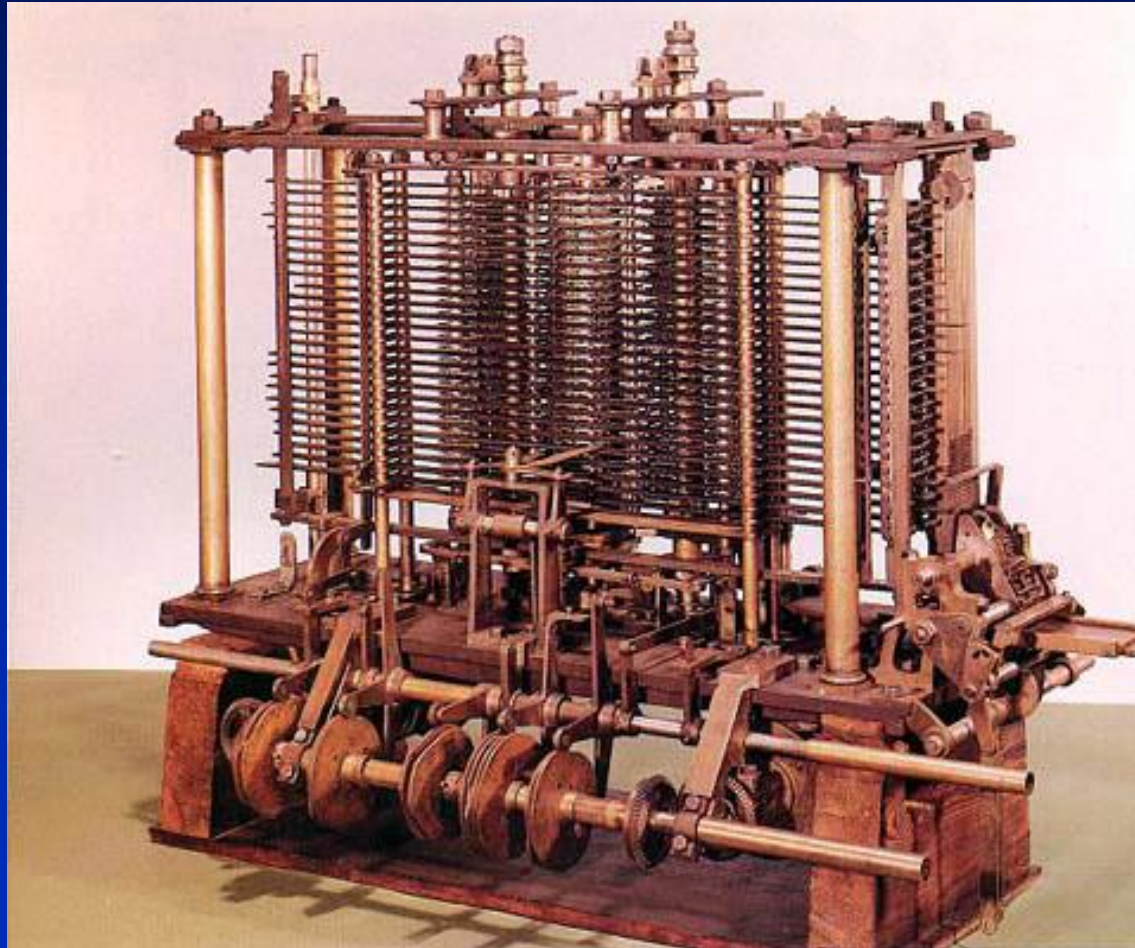
---

- *Required Reading:*

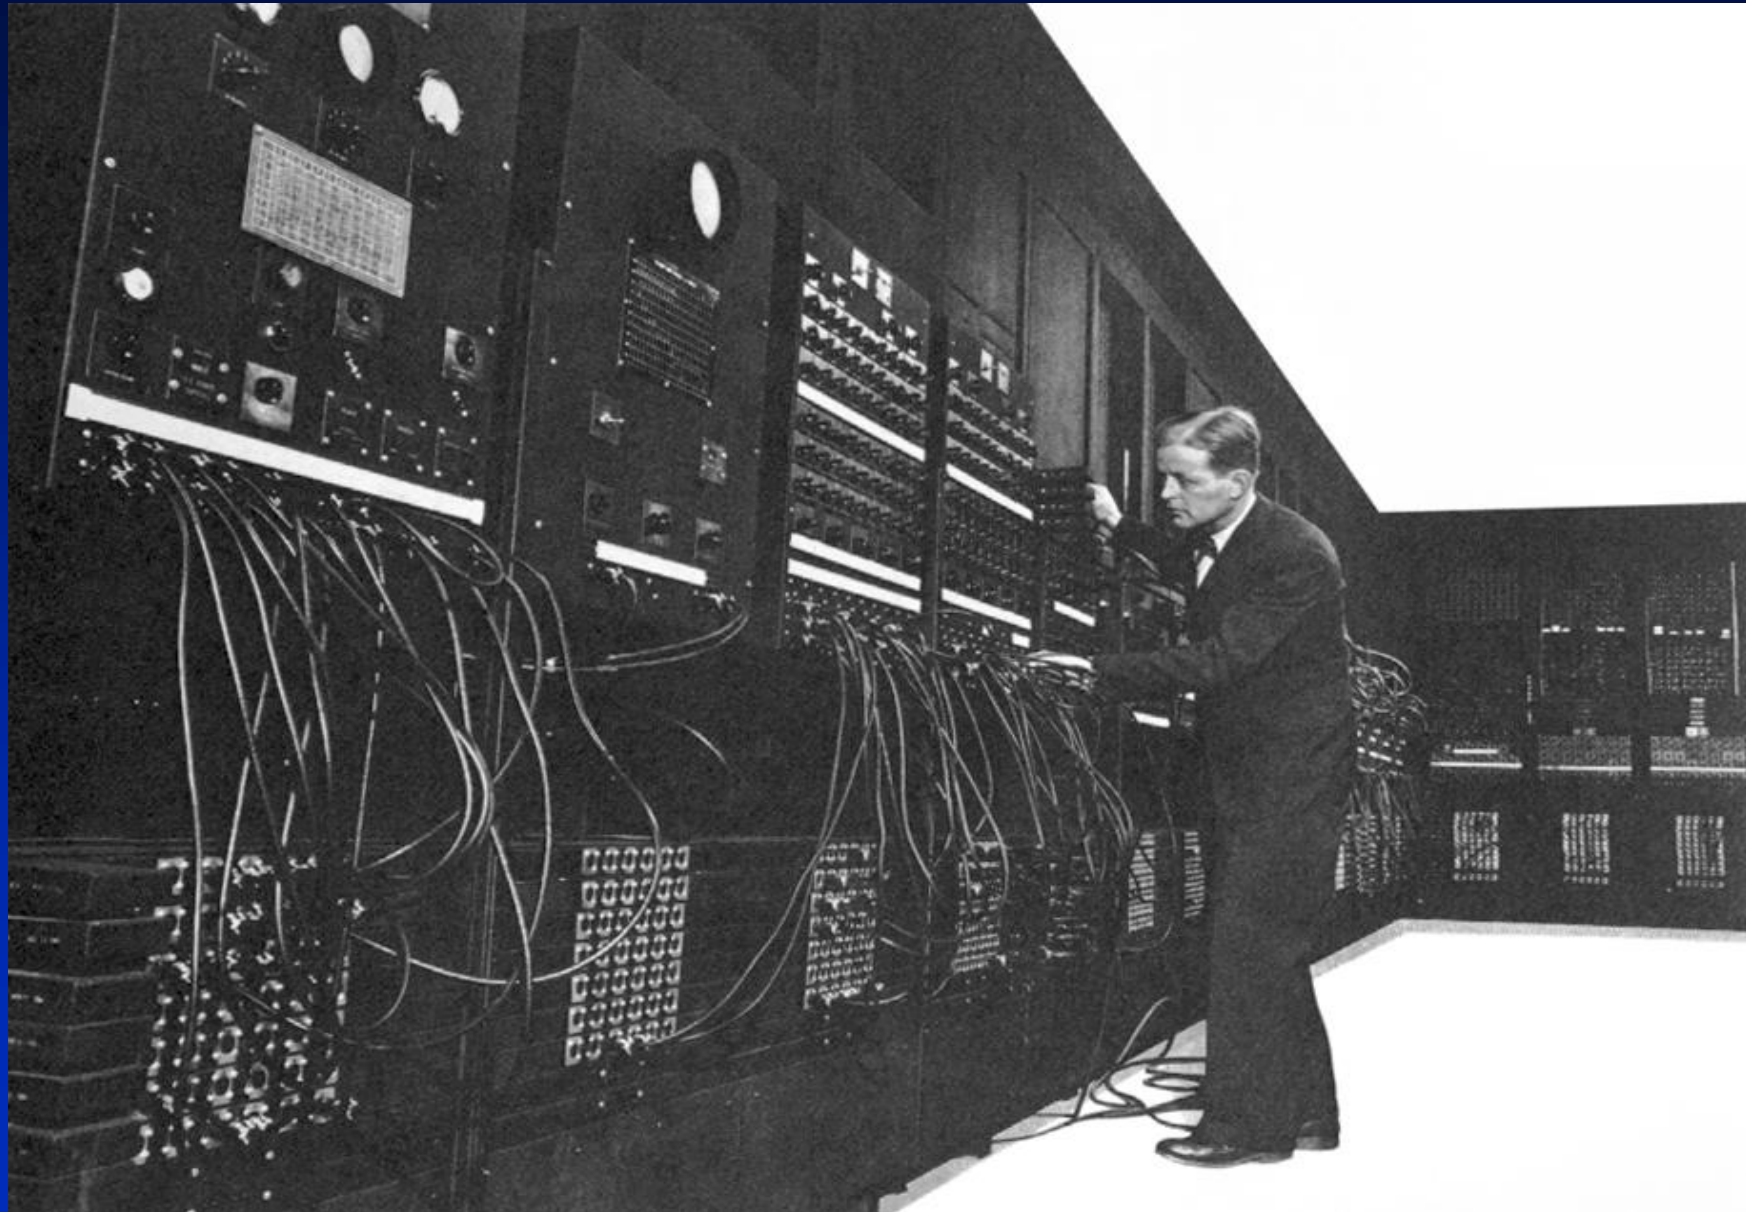
- Craig R. Barrett. From Sand to Silicon: Manufacturing an Integrated Circuit, *Scientific American*, Special Issue, The Solid-State Century, January 1998, pp. 55-61. (Search: e-Journals/ Scientific American Archive Online/article (full text) <http://www.library.cornell.edu/johnson/library/general/emba.html>)
- Mack, Chris. "The Multiple Lives of Moore's Law." *IEEE Spectrum* Apr. 2015: 30-37. *Cornell University Library*. Web. <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7065415>

# Turing Machine

# Alan Turing



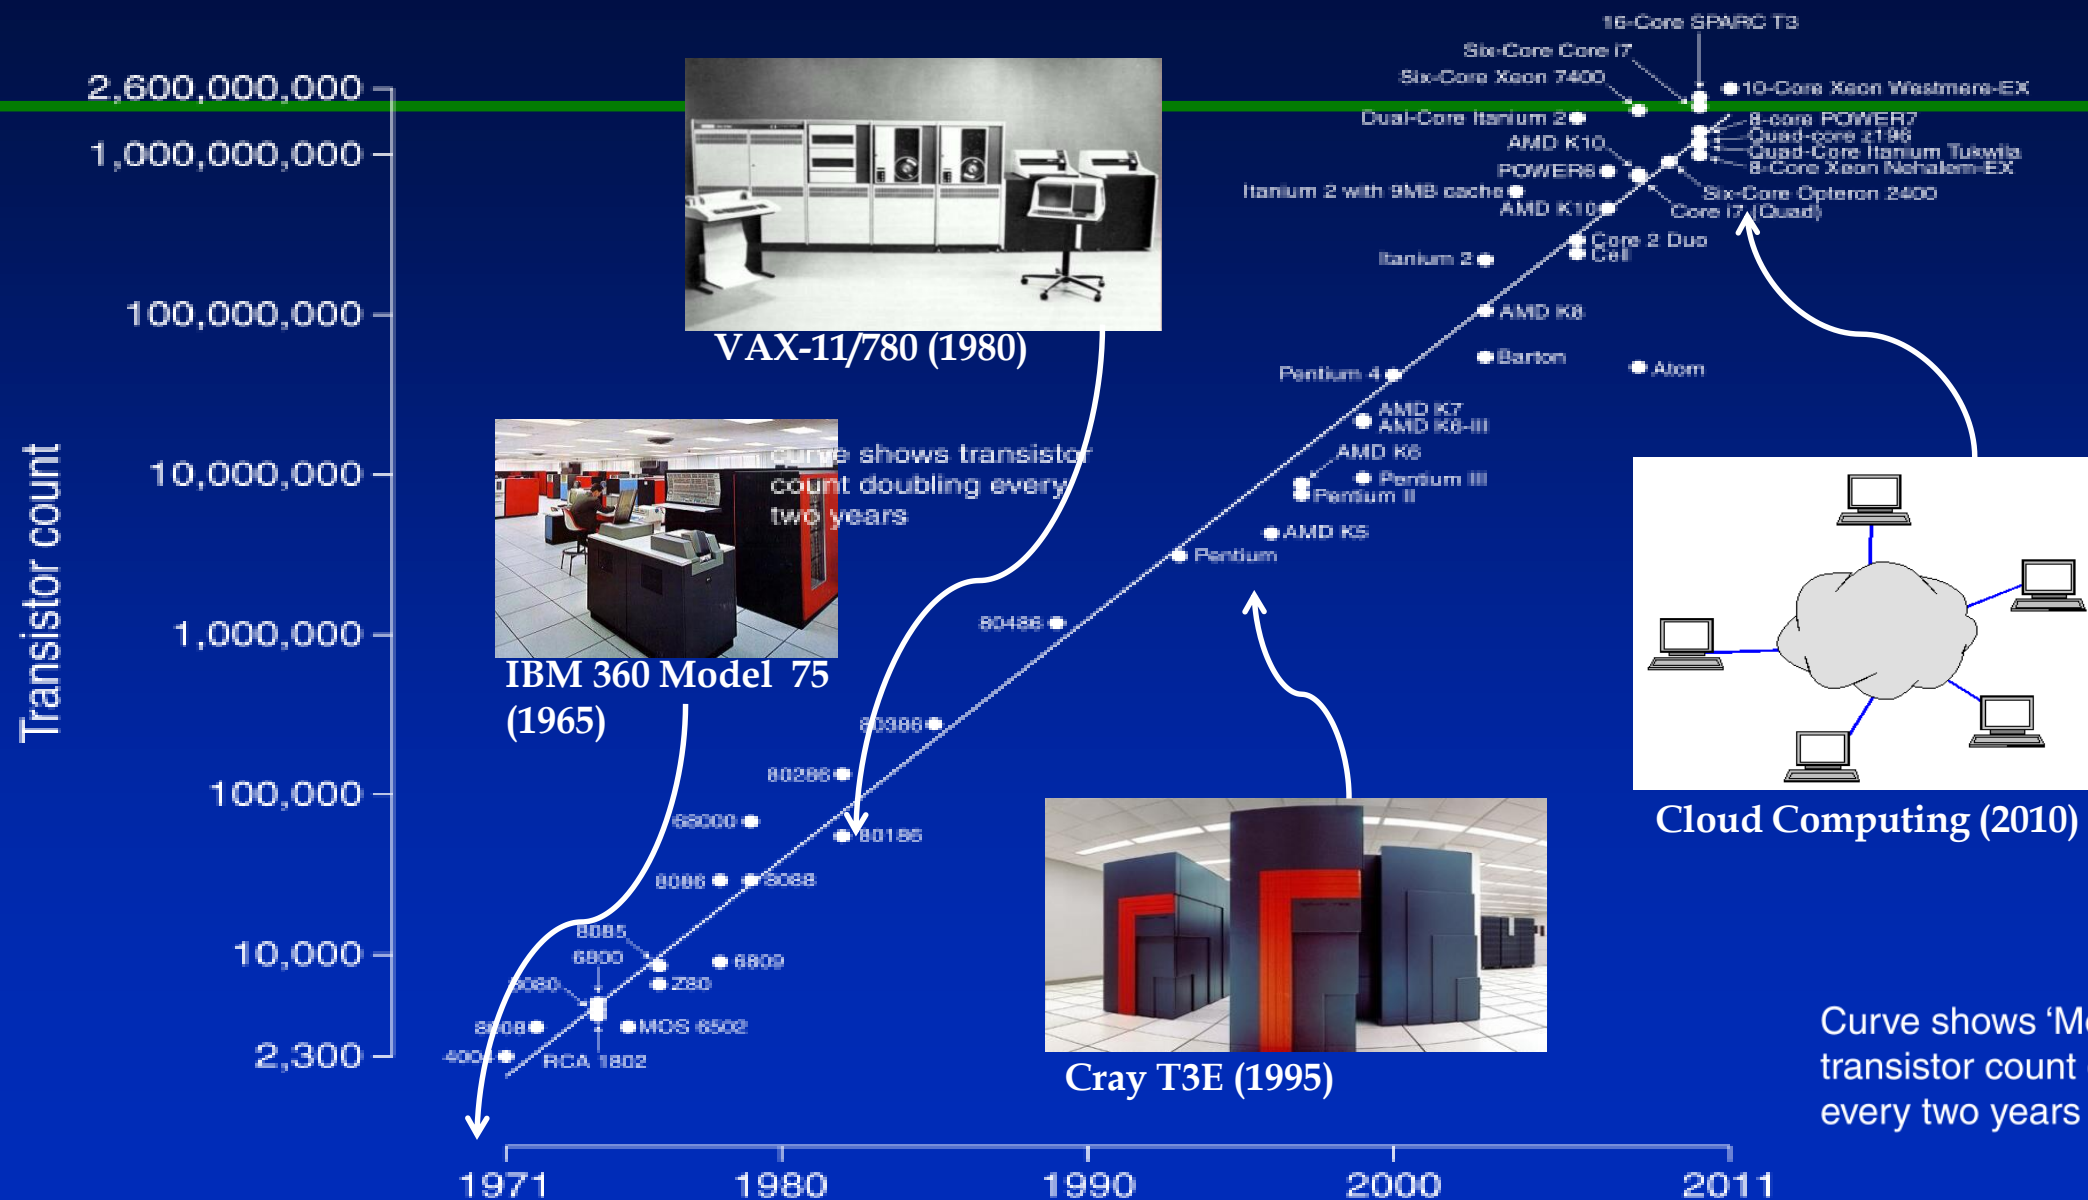
# Eniac 1946



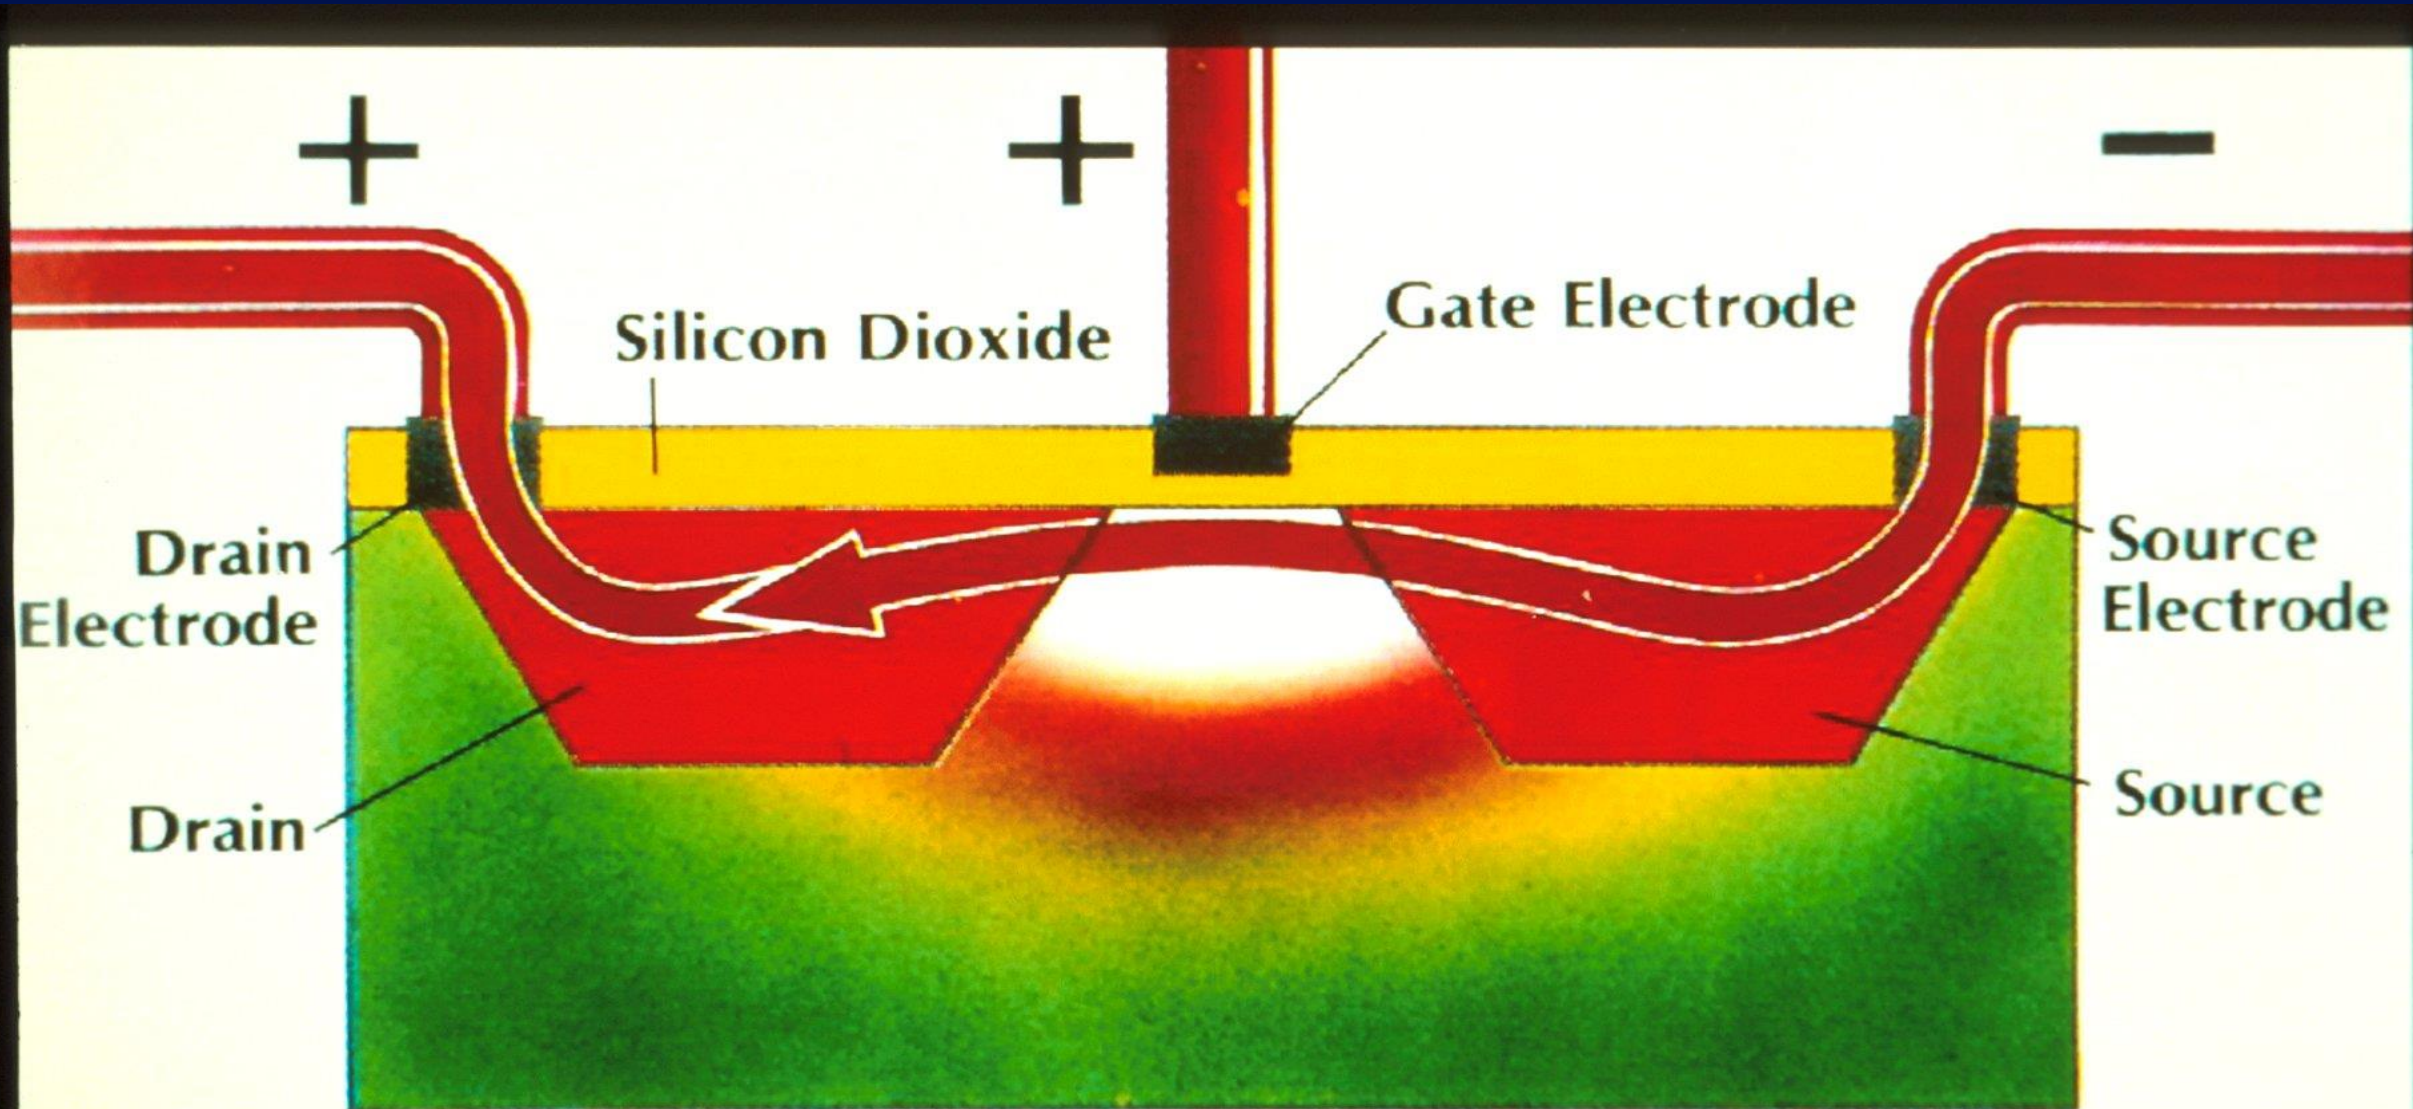
# Cloud Computing - 2010



# Microprocessor Transistor Counts 1971-2011 & Moore's Law

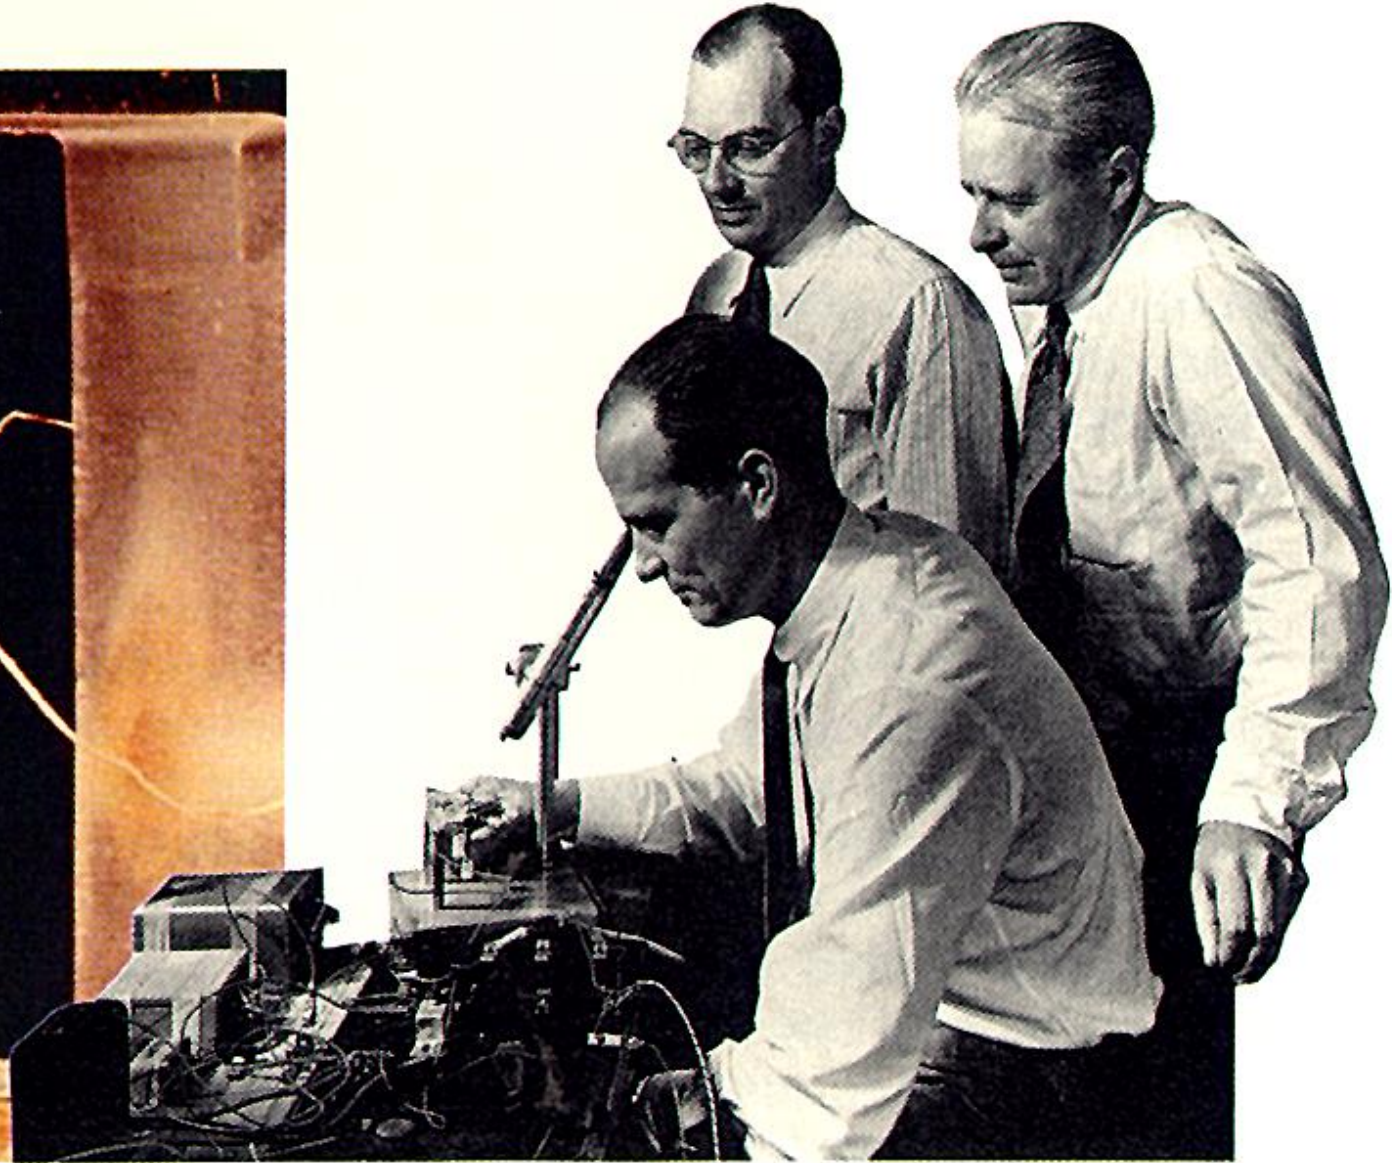
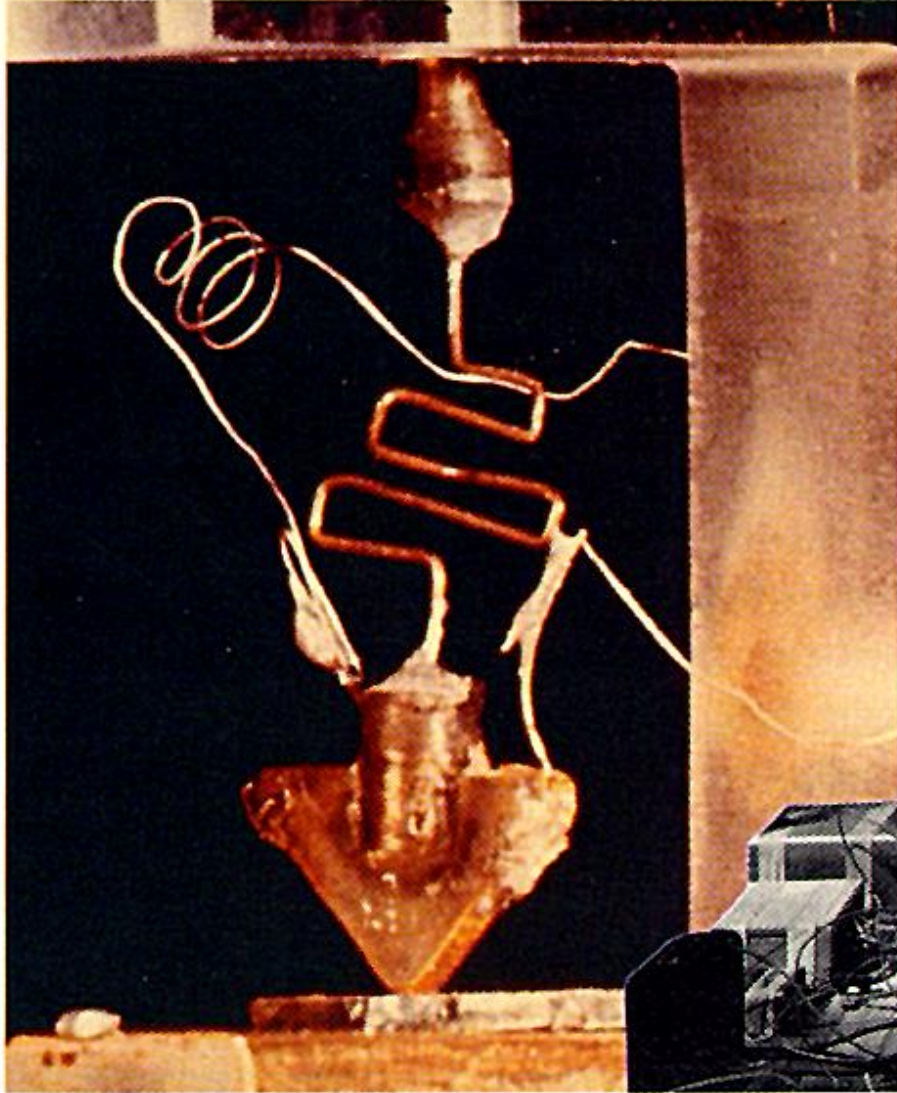


# Transistor



# Shockley, Bardeen & Brattain

1947



# Shockley & the Traitorous Eight

---

- William Shockley - Receives the Nobel Prize in Physics with Bardeen and Brattain (1956)  
leaves Bell Laboratory and forms Fairchild Semiconductor
- Julius Blank - founded Xicor
- Jean Hoerni - invented the planar process  
founded Amelco → Teledyne
- Jay Last - founded Amelco → Teledyne
- Sheldon Roberts - founded Amelco → Teledyne

# Shockley & the Traitorous Eight

---

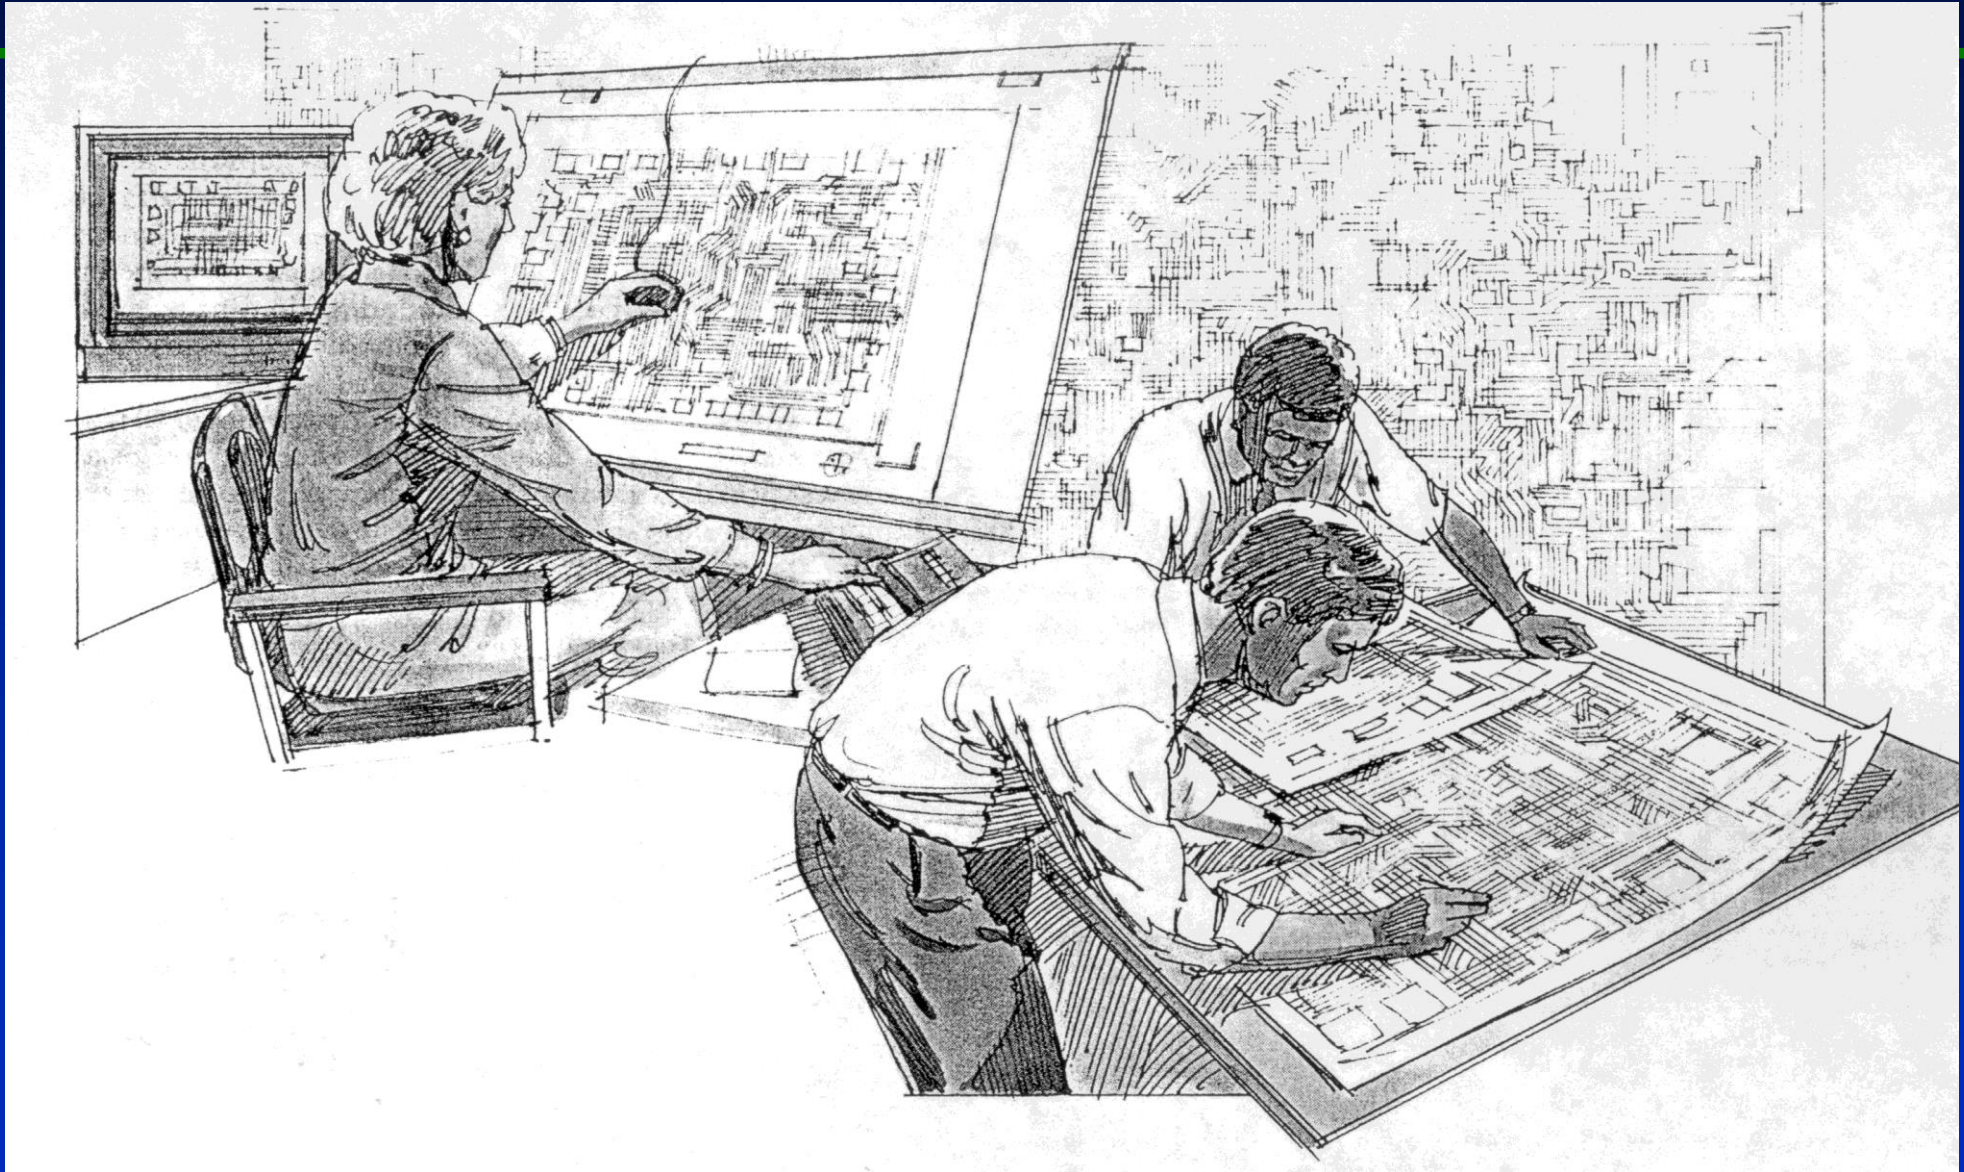
- Gordon Moore - founded Intel in 1968
- Robert Noyce - founded Intel in 1968
- Eugene Kleiner - founded Kleiner-Perkins
- Victor Grinich - only a poor professor at UC Berkeley & Stanford

---

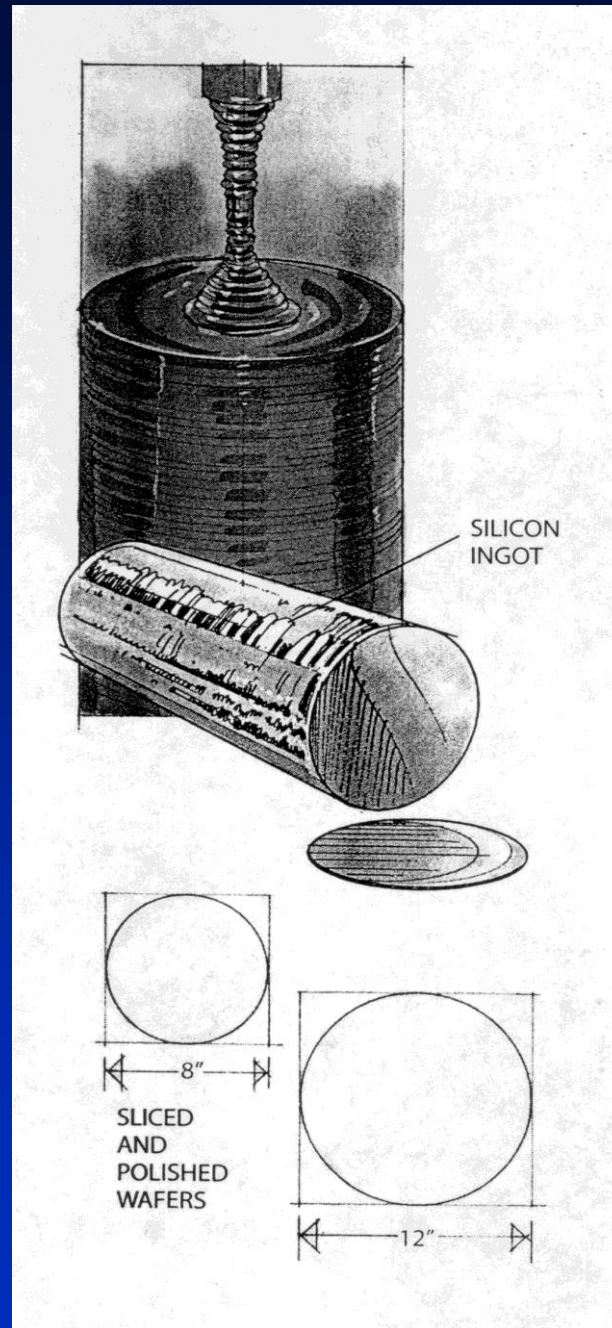
# From Sand to Silicon – Manufacturing an Integrated Circuit

Scientific American: The Solid-State  
Century, Special Issue 1998

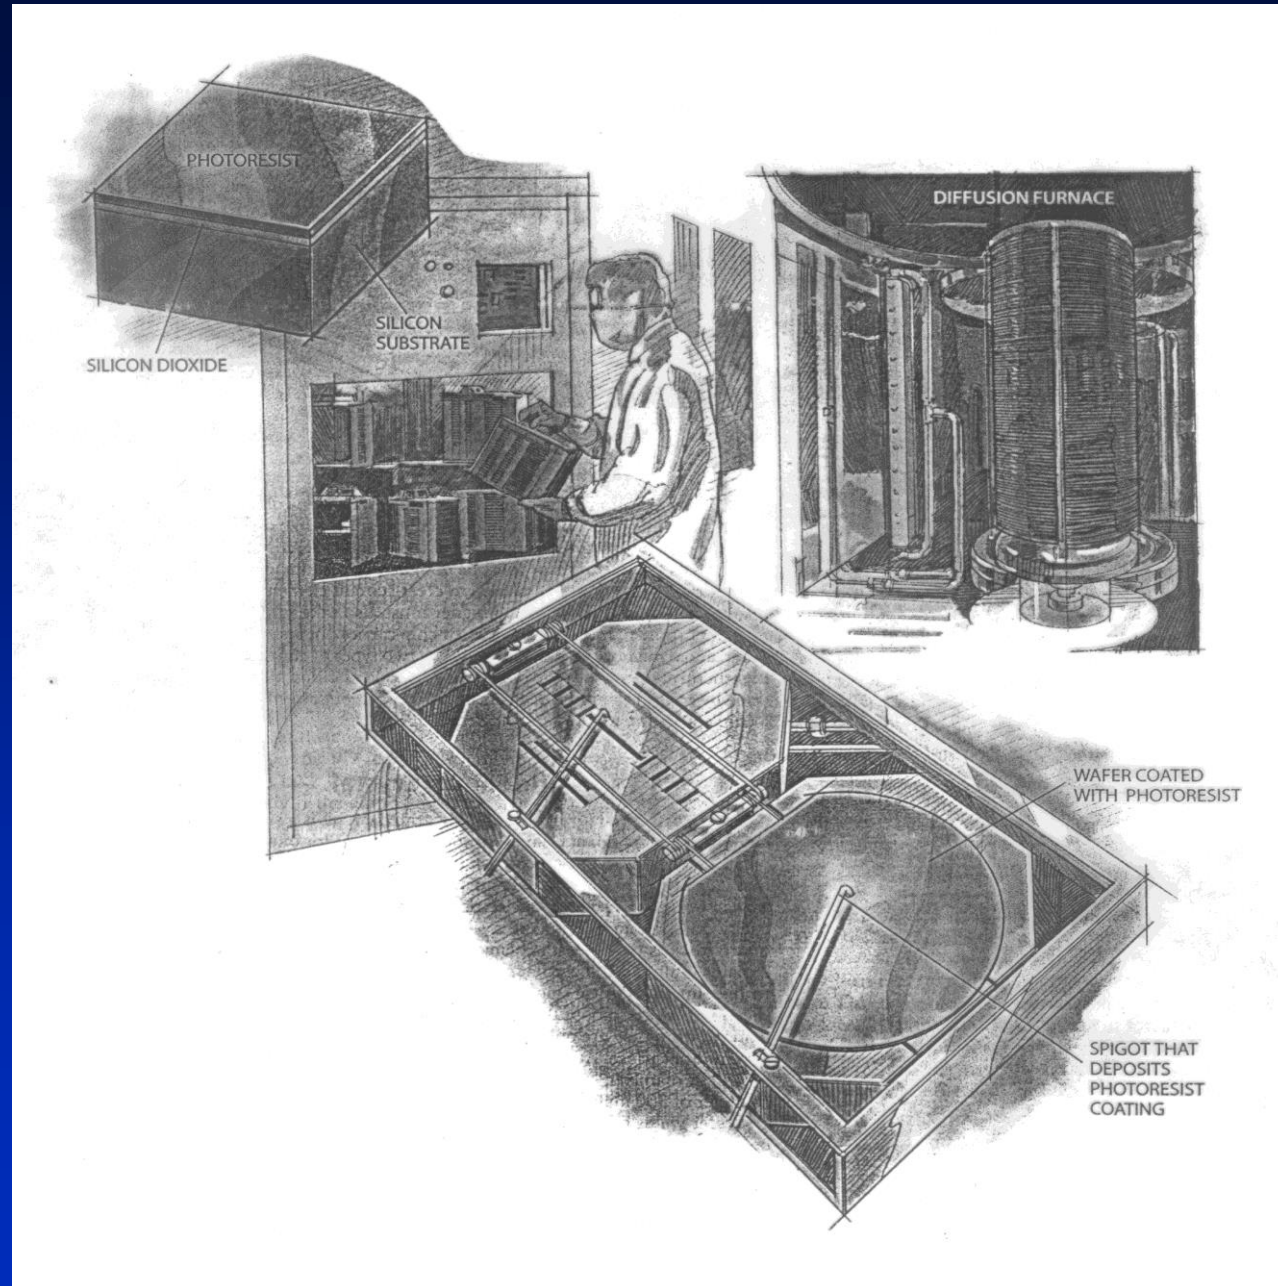
# Chip Design



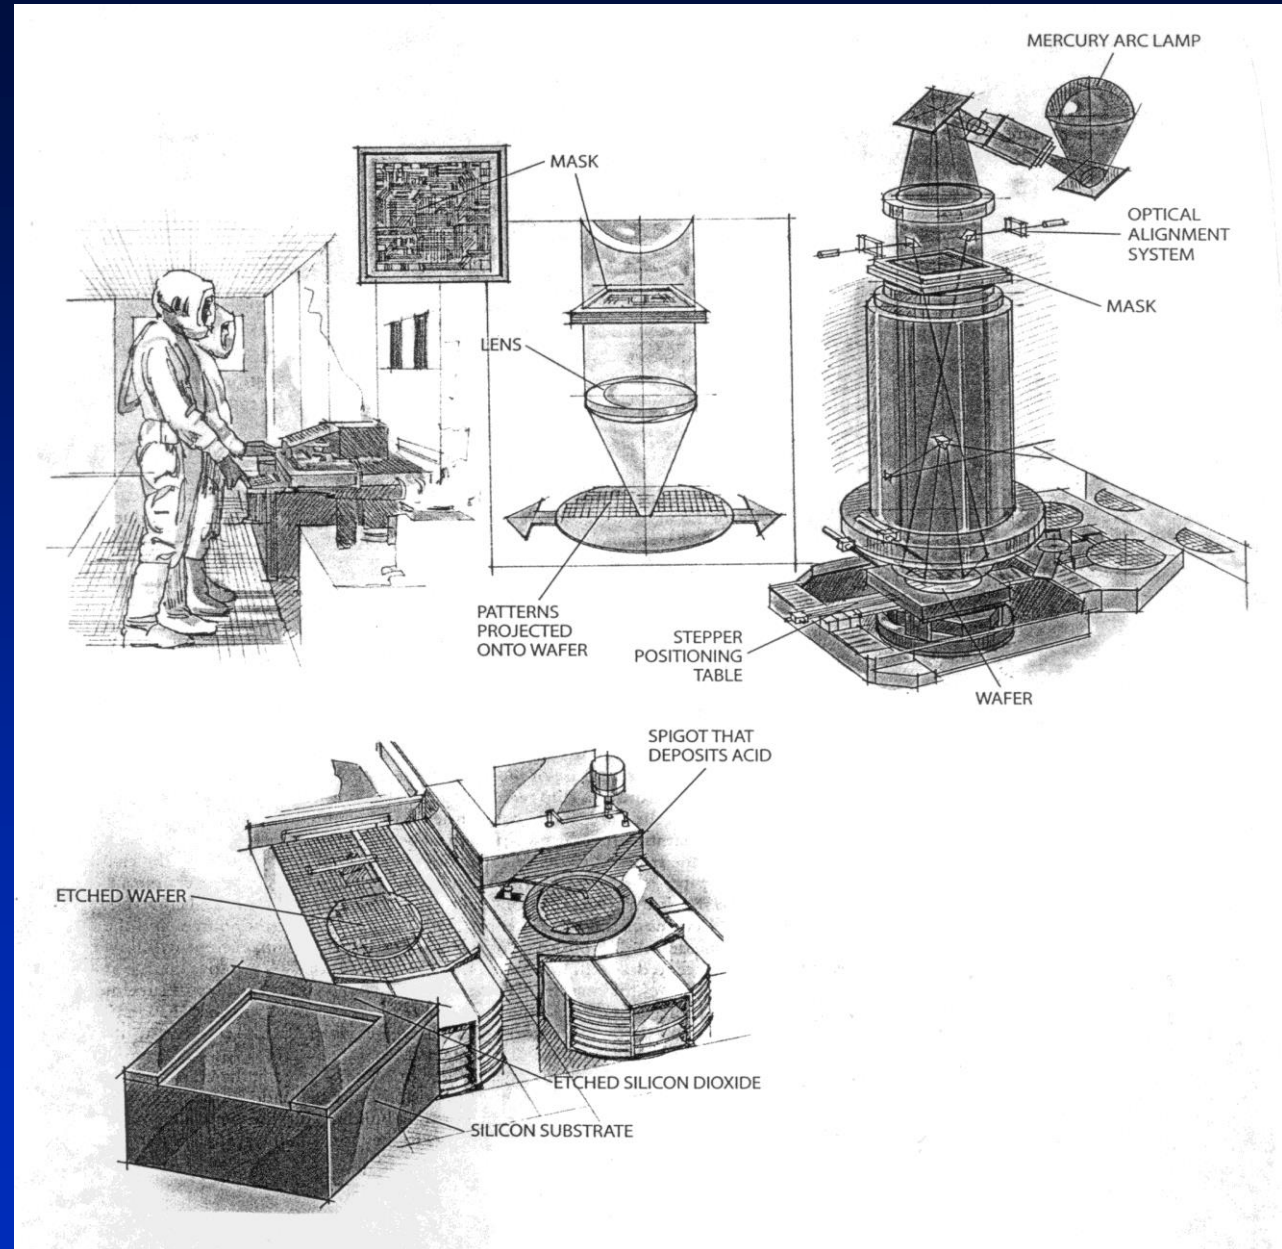
# Silicon Crystal



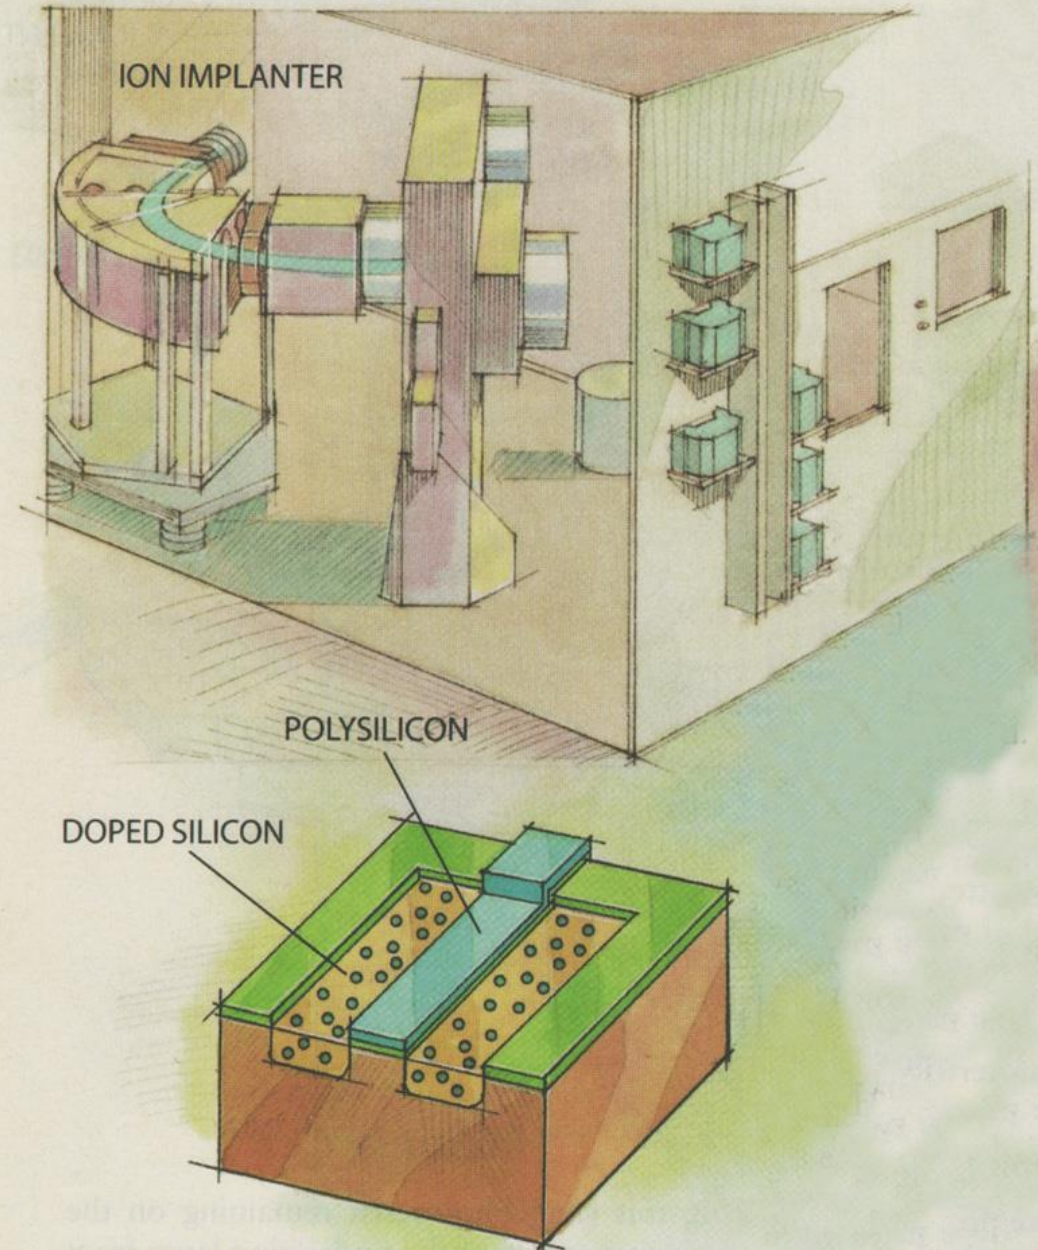
# Layering



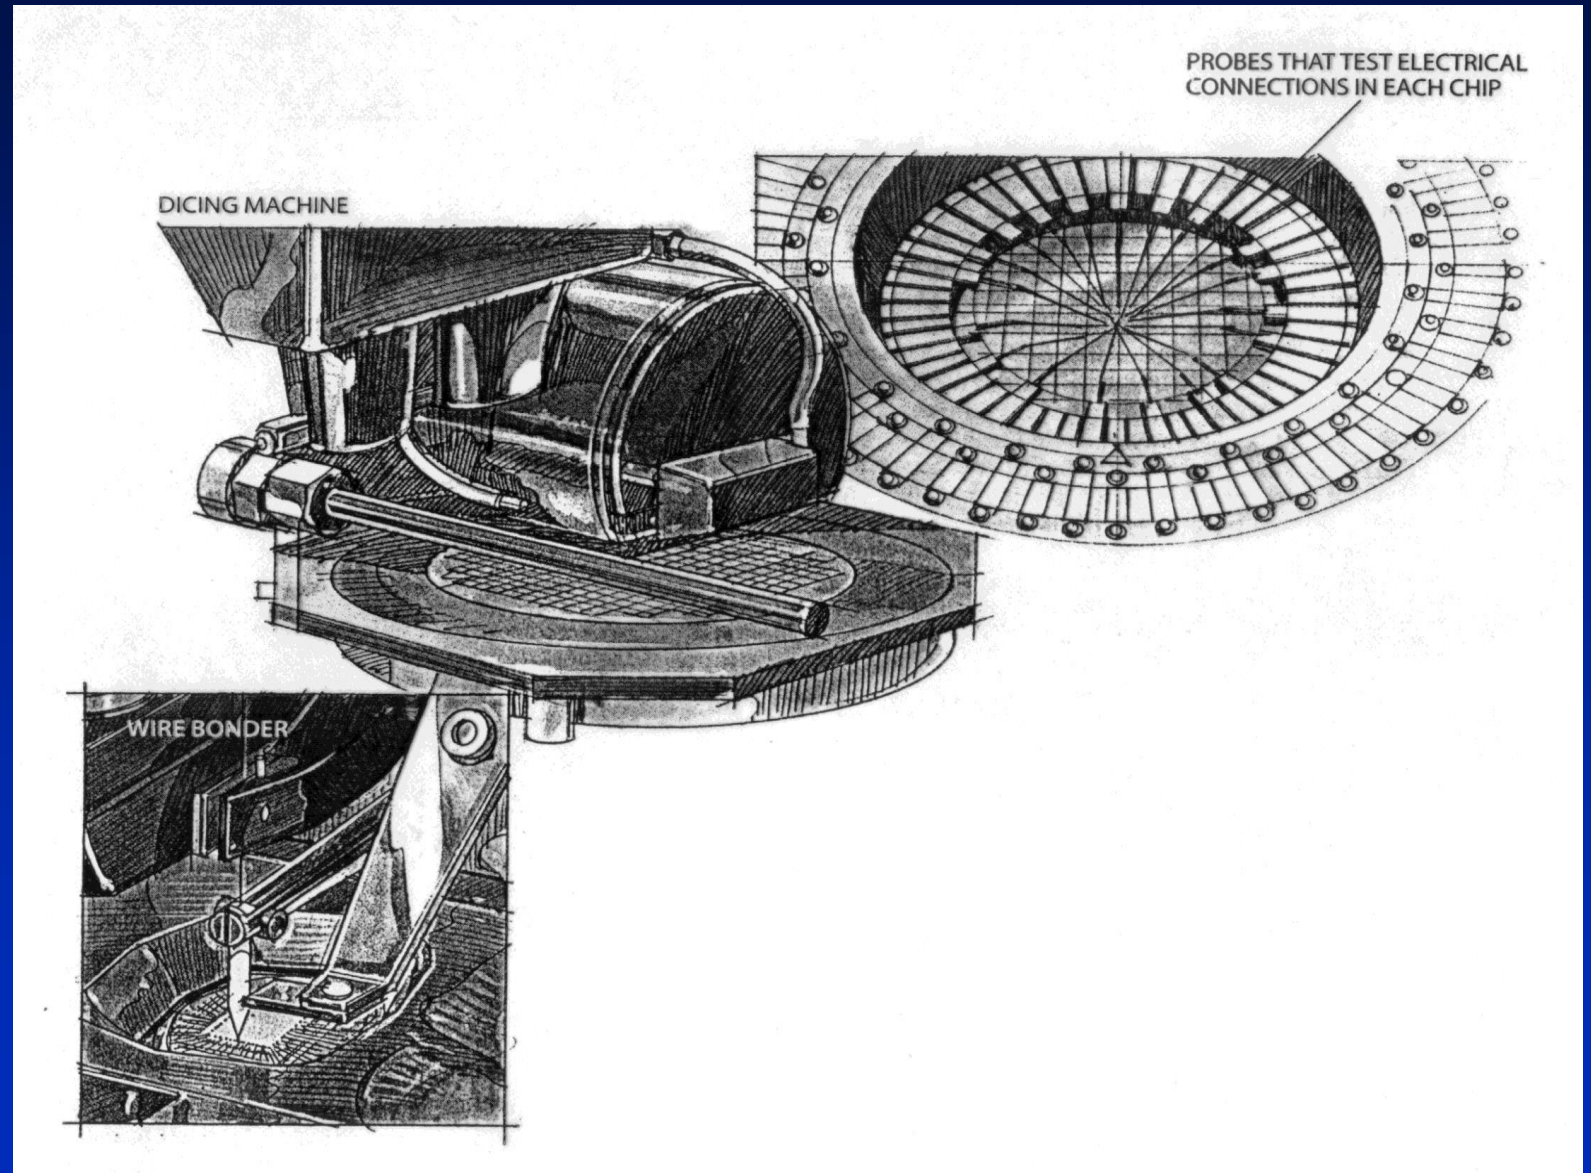
# Masking & Etching



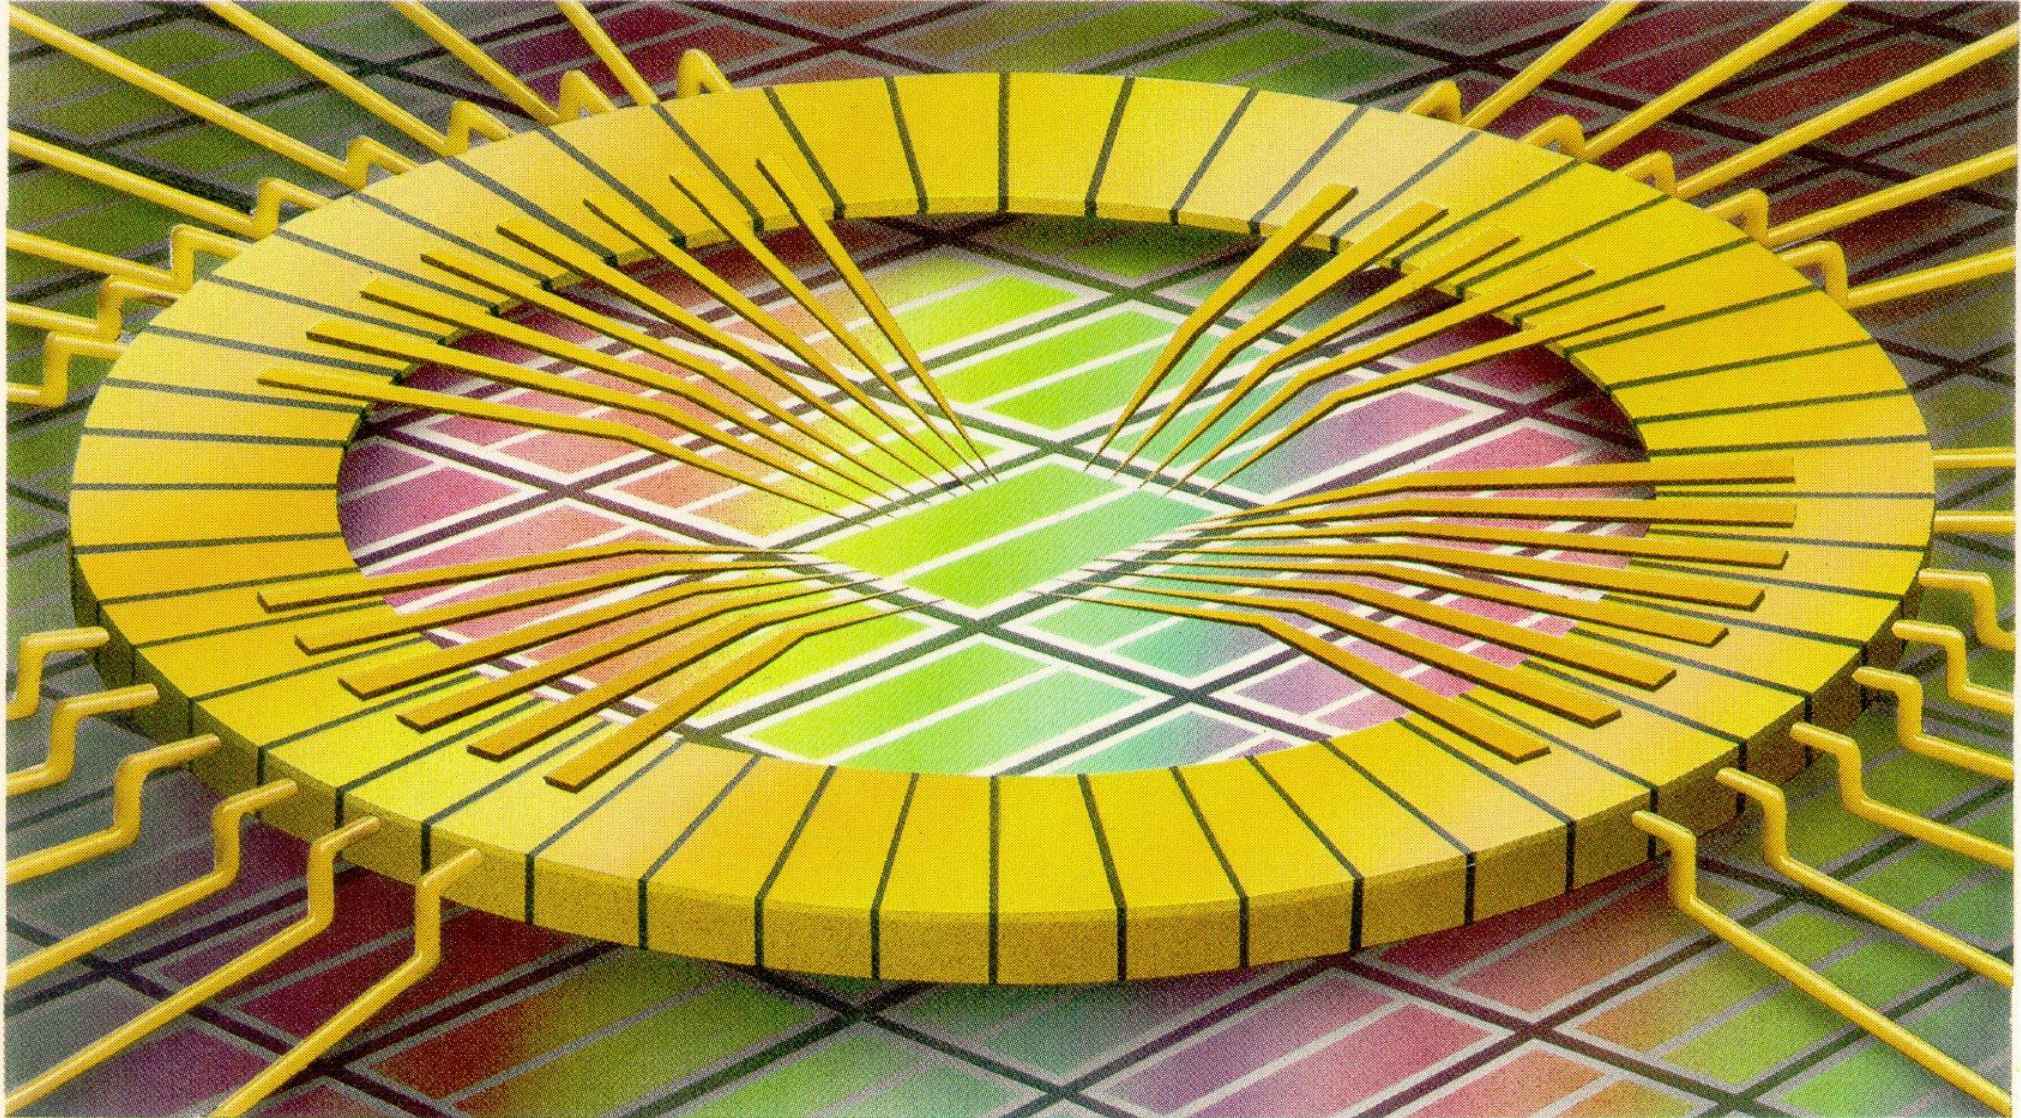
# Doping



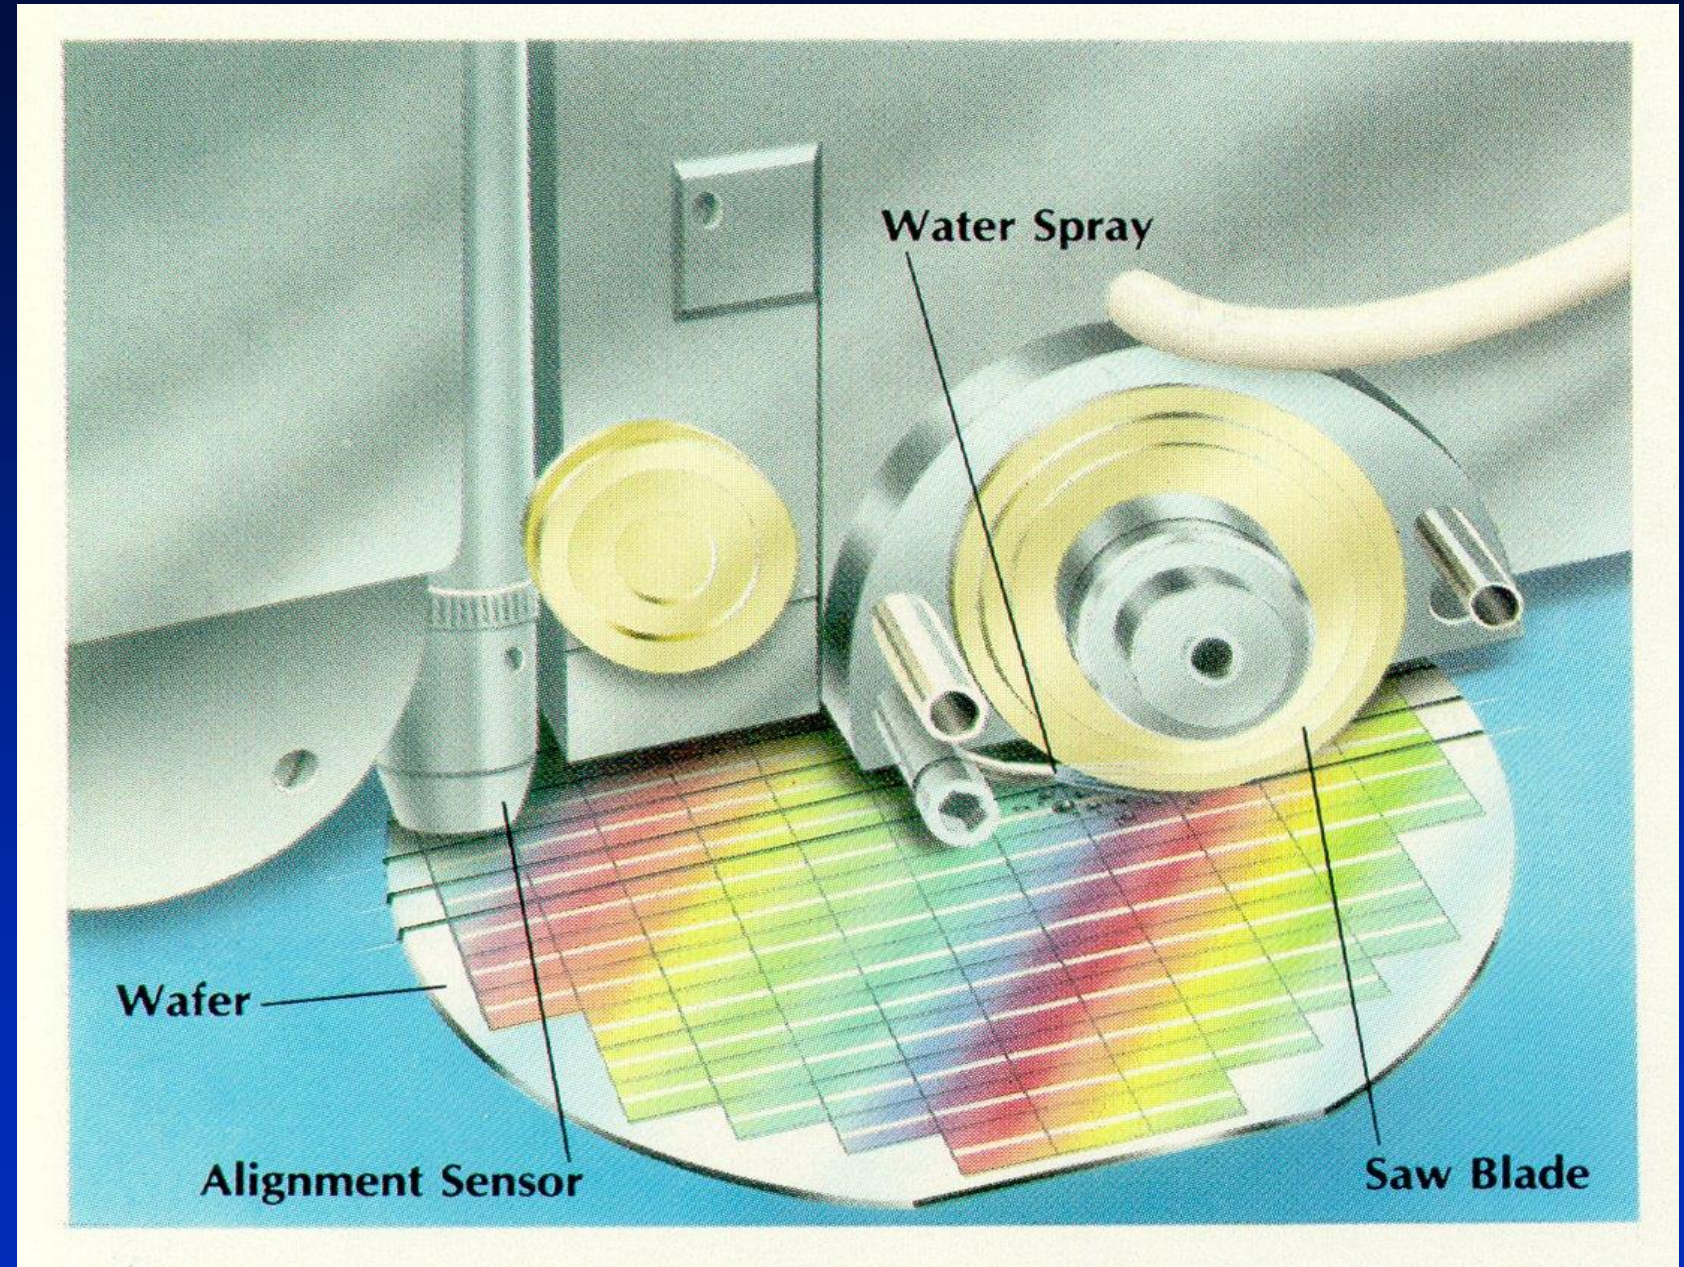
# Interconnections & Dicing



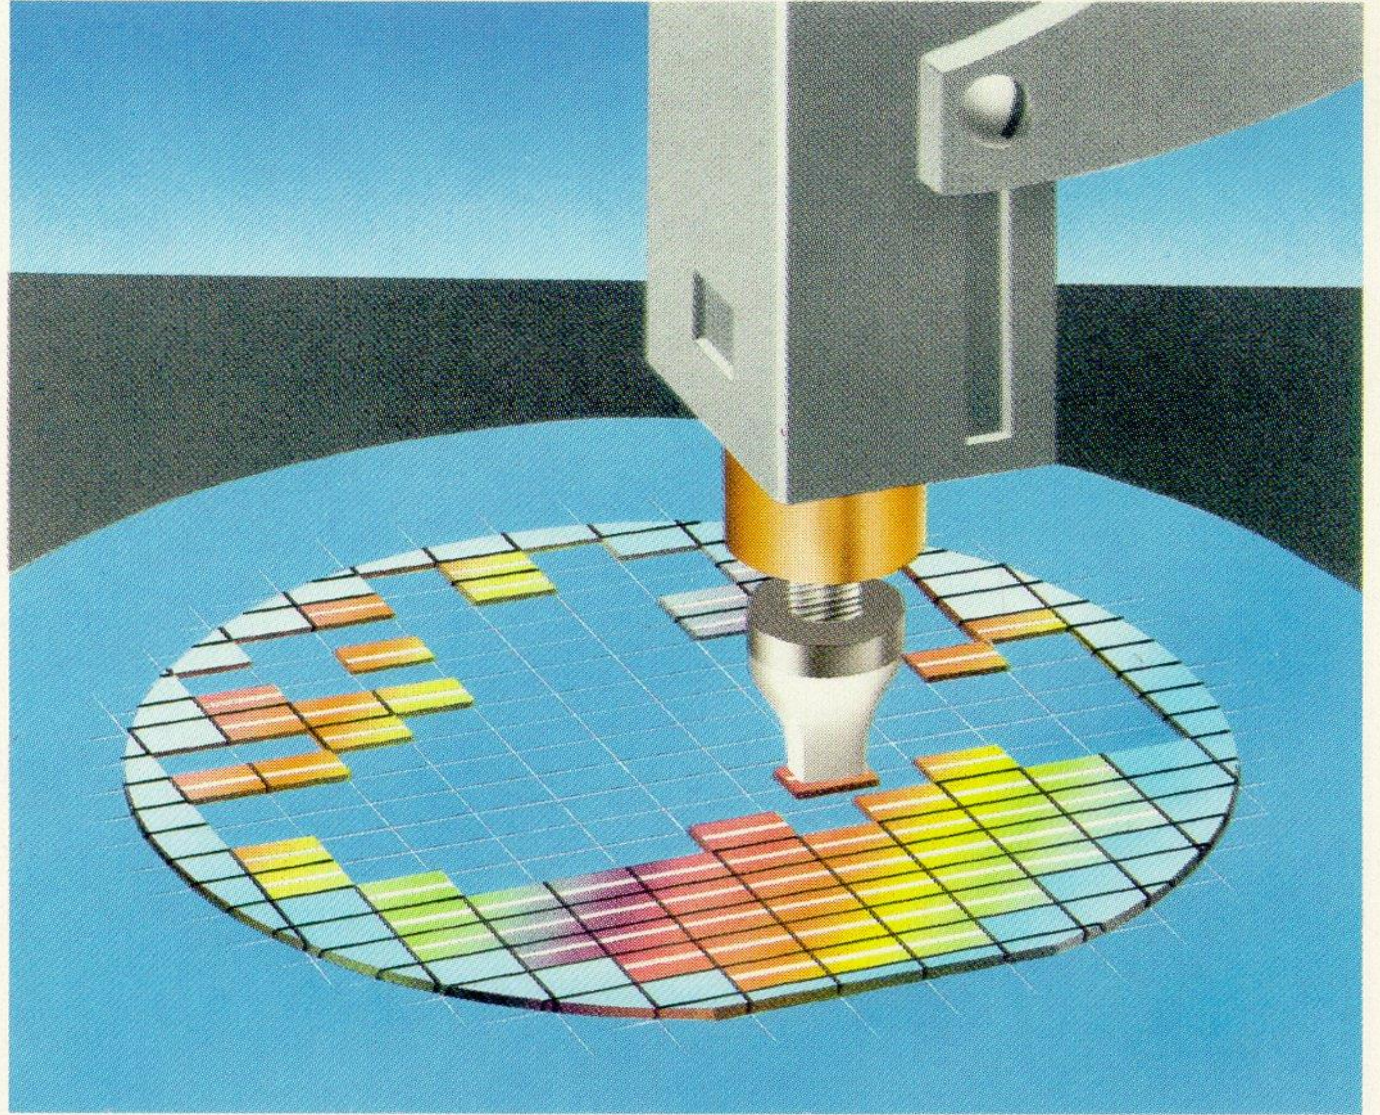
# Probing Electrical Connections



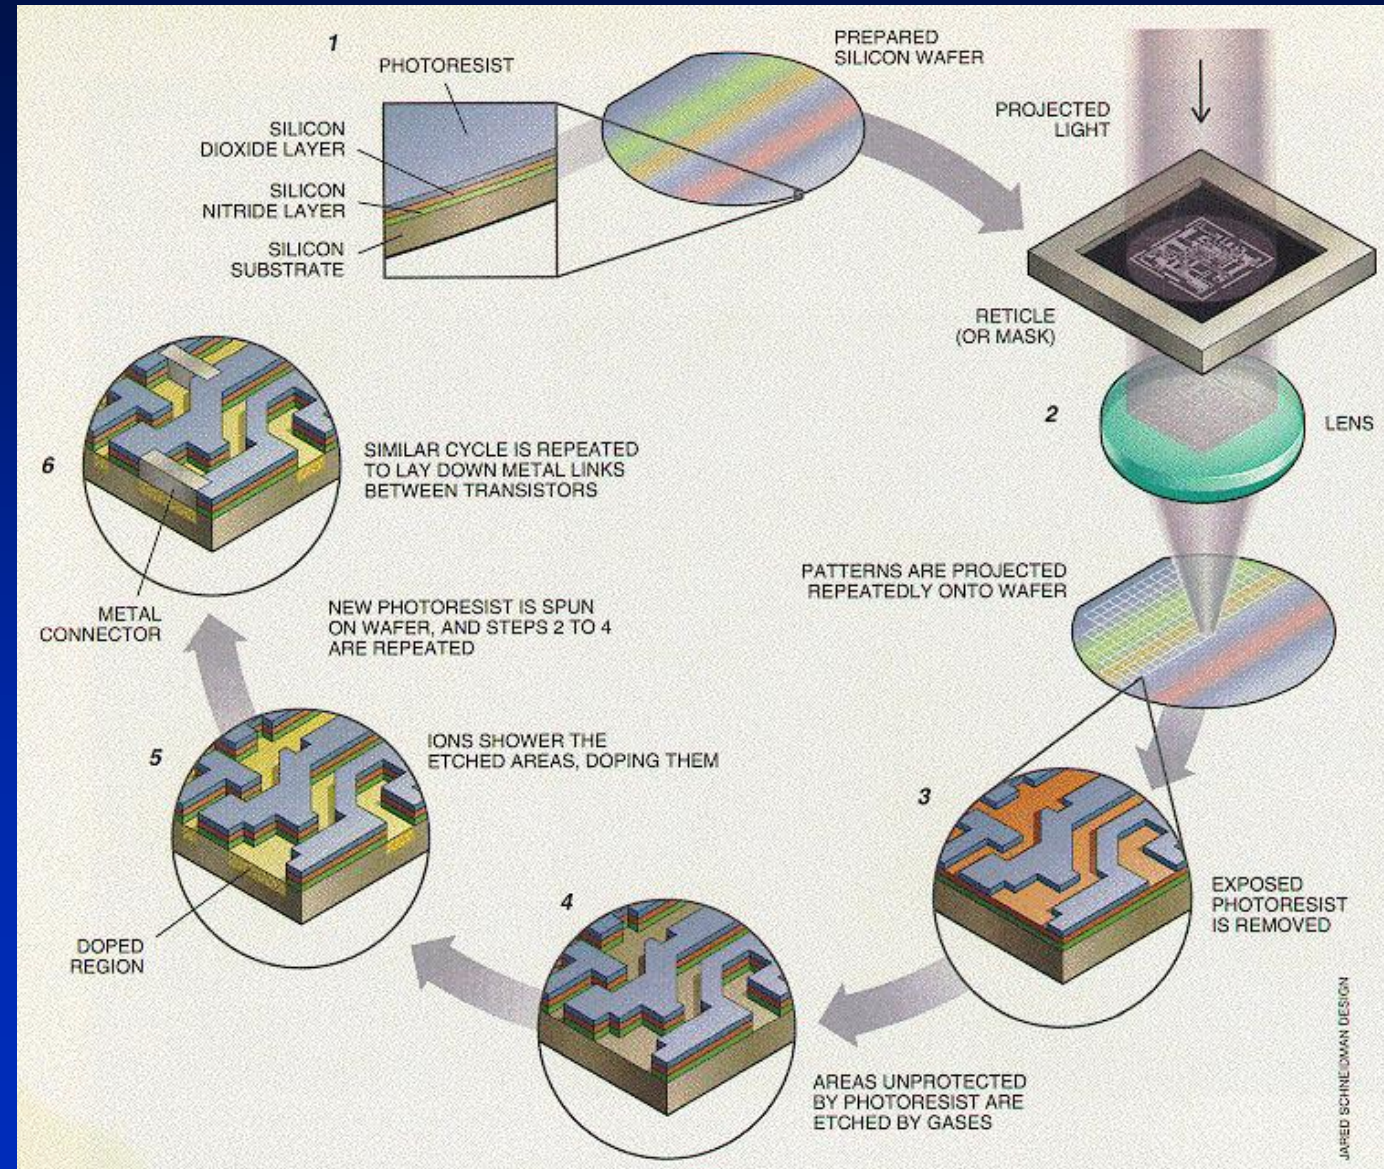
# Dicing

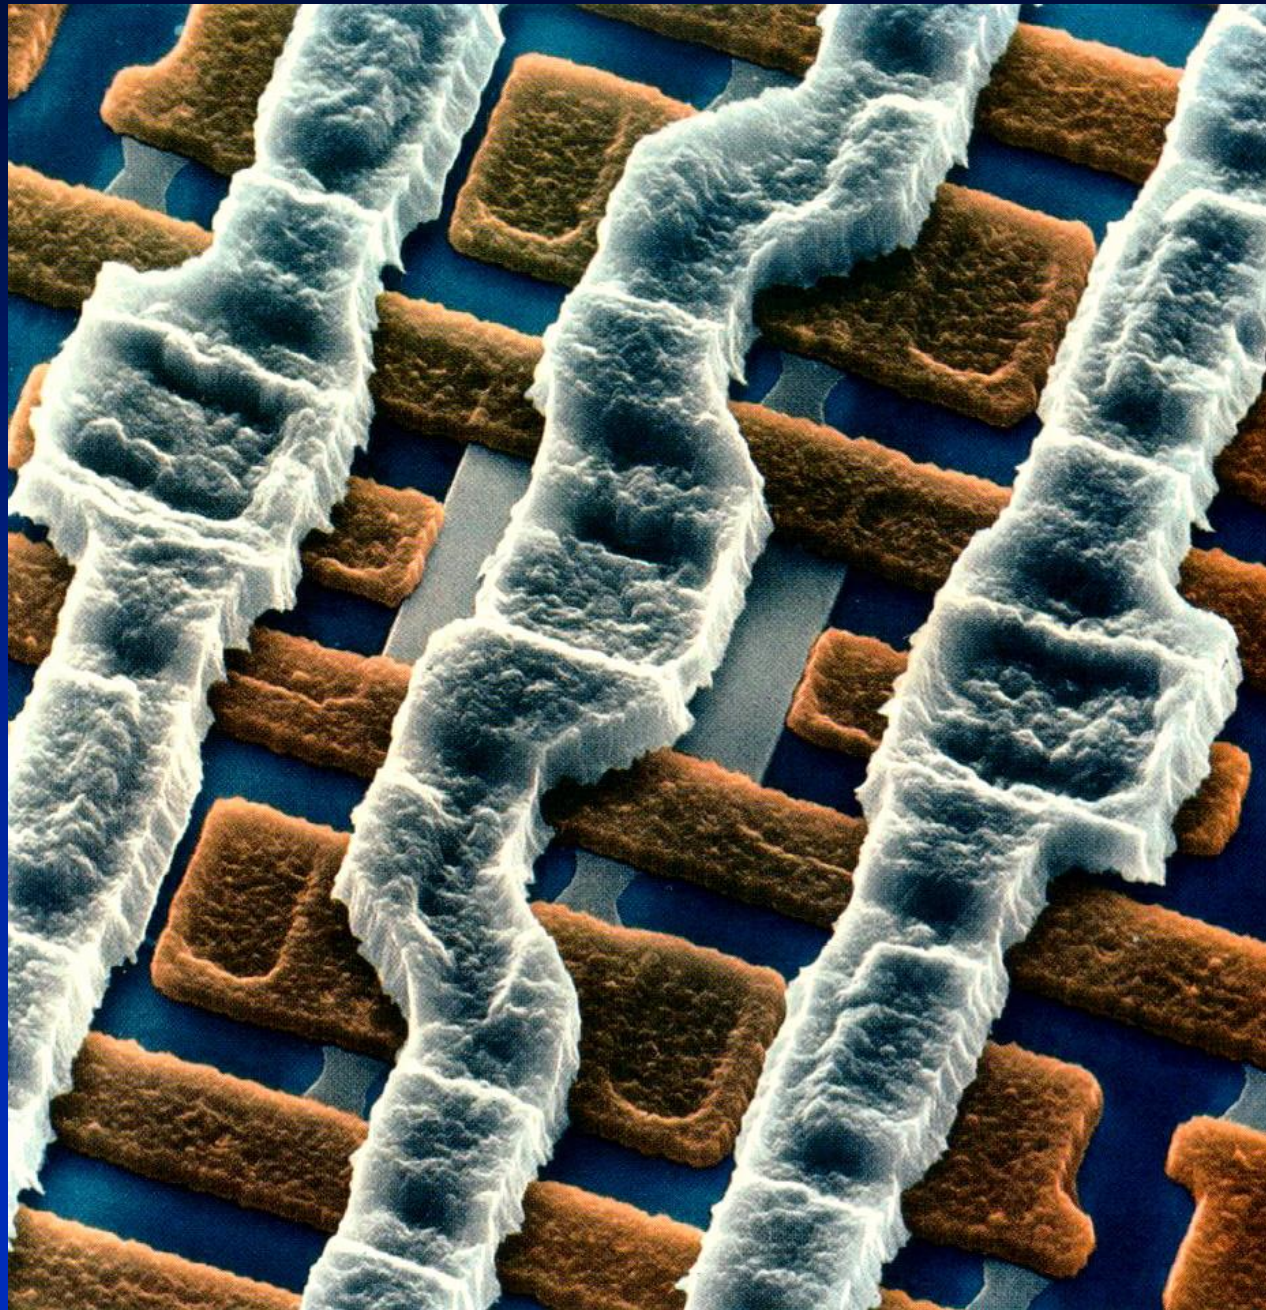


# Chip Selection



# Chip Fabrication





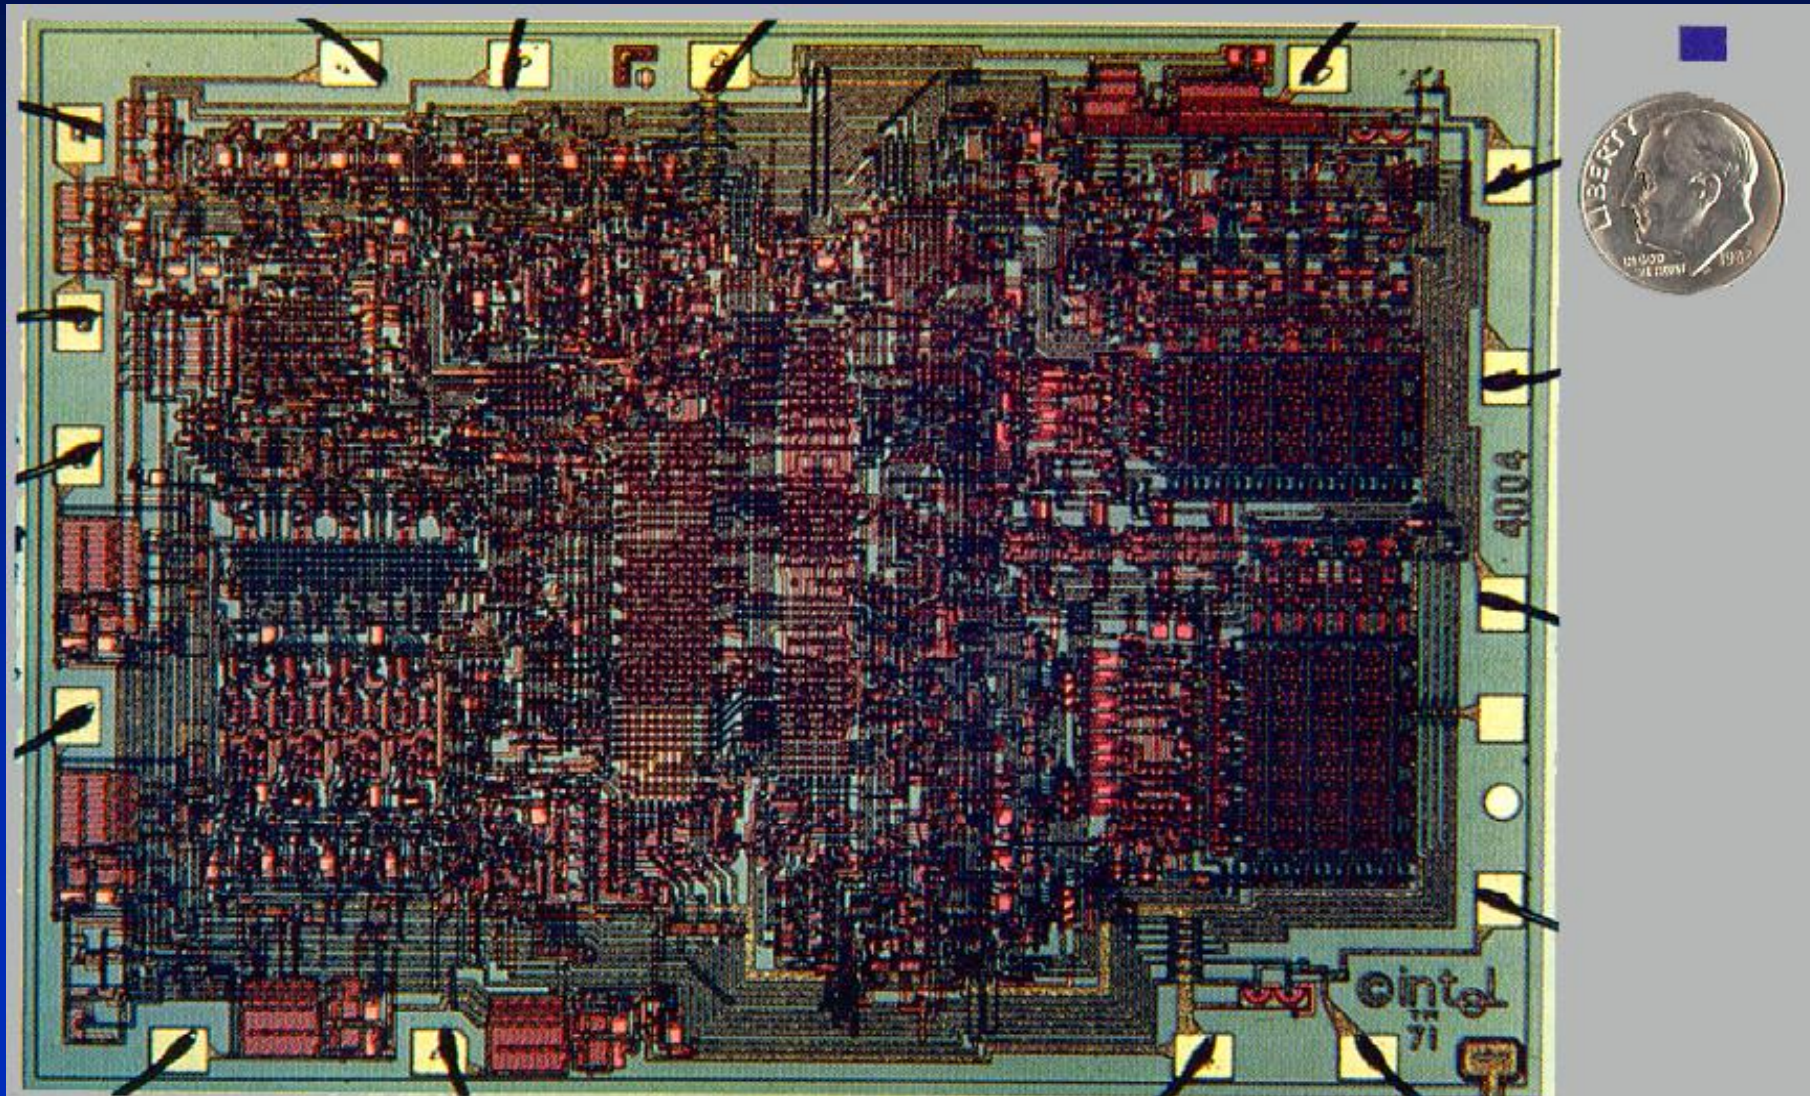
# International Technology Roadmap for Semiconductors

|                                     | 2001                | 2004                | 2007                | 2010                | 2013                | 2016                |
|-------------------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Technology<br>(nanometers)          | 130nm               | 90nm                | 65nm                | 45nm                | 32nm                | 22nm                |
| Functions<br>per Chip<br>(millions) | 97                  | 193                 | 386                 | 1546                | 3092                | 6184                |
| Clock Speed<br>(Ghz)                | 2.5Ghz              | 4.1Ghz              | 9.3Ghz              | 15Ghz               | 23Ghz               | 40Ghz               |
| Wafer Size<br>(millimeters)         | 200mm               | 300mm               | 300mm               | 300mm               | 450mm               | 450mm               |
| Chip Size<br>(mm <sup>2</sup> )     | 140 mm <sup>2</sup> | 140 mm <sup>2</sup> | 140 mm <sup>2</sup> | 140 mm <sup>2</sup> | 140 mm <sup>2</sup> | 140 mm <sup>2</sup> |

Roughly 0.5 shrink every 3 years. Intel released 22 nm chips in 2013

# Intel 4004

# November 1971



# Moore's Law – CPU Transistor Counts

| Processor   | Transistor count | Date of introduction | Manufacturer | Process | Area |
|-------------|------------------|----------------------|--------------|---------|------|
| Intel 4004  | 2,300            | 1971                 | Intel        | 10 µm   |      |
| Intel 8008  | 3,500            | 1972                 | Intel        | 10 µm   |      |
| Intel 8080  | 4,500            | 1974                 | Intel        | 6 µm    |      |
| Intel 8088  | 29,000           | 1979                 | Intel        | 3 µm    |      |
| Intel 80286 | 134,000          | 1982                 | Intel        | 1.5 µm  |      |
| Intel 80386 | 275,000          | 1985                 | Intel        | 1.5 µm  |      |
| Intel 80486 | 1,180,000        | 1989                 | Intel        | 1 µm    |      |
| Pentium     | 3,100,000        | 1993                 | Intel        | 0.8 µm  |      |
| AMD K5      | 4,300,000        | 1996                 | AMD          | 0.5 µm  |      |
| Pentium II  | 7,500,000        | 1997                 | Intel        | 0.35 µm |      |
| AMD K6      | 8,800,000        | 1997                 | AMD          | 0.35 µm |      |
| Pentium III | 9,500,000        | 1999                 | Intel        | 0.25 µm |      |
| AMD K6-III  | 21,300,000       | 1999                 | AMD          | 0.25 µm |      |
| AMD K7      | 22,000,000       | 1999                 | AMD          | 0.25 µm |      |
| Pentium 4   | 42,000,000       | 2000                 | Intel        | 180 nm  |      |
| Atom        | 47,000,000       | 2008                 | Intel        | 45 nm   |      |
| Barton      | 54,300,000       | 2003                 | AMD          | 130 nm  |      |
| AMD K8      | 105,900,000      | 2003                 | AMD          | 130 nm  |      |
| Itanium 2   | 220,000,000      | 2003                 | Intel        | 130 nm  |      |

# Moore's Law – CPU Transistor Counts

| Processor                         | Transistor count | Date of introduction | Manufacturer | Process | Area                |
|-----------------------------------|------------------|----------------------|--------------|---------|---------------------|
| Core 2 Duo                        | 291,000,000      | 2006                 | Intel        | 65 nm   |                     |
| AMD K10                           | 463,000,000      | 2007                 | AMD          | 65 nm   |                     |
| AMD K10                           | 758,000,000      | 2008                 | AMD          | 45 nm   |                     |
| Itanium 2 with 9MB cache          | 592,000,000      | 2004                 | Intel        | 130 nm  |                     |
| Core i7 (Quad)                    | 731,000,000      | 2008                 | Intel        | 45 nm   | 263 mm <sup>2</sup> |
| POWER6                            | 789,000,000      | 2007                 | IBM          | 65 nm   | 341 mm <sup>2</sup> |
| Six-Core Opteron 2400             | 904,000,000      | 2009                 | AMD          | 45 nm   |                     |
| Six-Core Core i7                  | 1,170,000,000    | 2010                 | Intel        | 32 nm   |                     |
| Dual-Core Itanium 2               | 1,700,000,000    | 2006                 | Intel        | 90 nm   | 596 mm <sup>2</sup> |
| Six-Core Xeon 7400                | 1,900,000,000    | 2008                 | Intel        | 45 nm   |                     |
| Quad-Core Itanium Tukwila         | 2,000,000,000    | 2010                 | Intel        | 65 nm   |                     |
| Six-Core Core i7 (Sandy Bridge-E) | 2,270,000,000    | 2011                 | Intel        | 32 nm   | 434 mm <sup>2</sup> |
| 8-Core Xeon Nehalem-EX            | 2,300,000,000    | 2010                 | Intel        | 45 nm   | 684 mm <sup>2</sup> |
| 10-Core Xeon Westmere-EX          | 2,600,000,000    | 2011                 | Intel        | 32 nm   | 512 mm <sup>2</sup> |
| Six-core zEC12                    | 2,750,000,000    | 2012                 | IBM          | 32 nm   | 597 mm <sup>2</sup> |
| 8-Core Itanium Poulson            | 3,100,000,000    | 2012                 | Intel        | 32 nm   | 544 mm <sup>2</sup> |
| 15-Core Xeon Ivy Bridge-EX        | 4,310,000,000    | 2014                 | Intel        | 22nm    | 541 mm <sup>2</sup> |
| 62-Core Xeon Phi                  | 5,000,000,000    | 2012                 | Intel        | 22 nm   | 350 mm <sup>2</sup> |
| Xbox One Main SoC                 | 5,000,000,000    | 2013                 | Microsoft    | 28 nm   | 363 mm <sup>2</sup> |
| 18-core Xeon Haswell-E5           | 5,560,000,000    | 2014                 | Intel        | 22 nm   | 661mm <sup>2</sup>  |
| IBM z13 Storage Controller        | 7,100,000,000    | 2015                 | IBM          | 22 nm   | 678 mm <sup>2</sup> |

# Moore's Law – GPU Transistor Counts

| Processor       | Transistor count | Date of introduction | Manufacturer | Process | Area                |
|-----------------|------------------|----------------------|--------------|---------|---------------------|
| R520            | 321,000,000      | 2005                 | AMD          | 90 nm   | 288 mm <sup>2</sup> |
| R580            | 384,000,000      | 2006                 | AMD          | 90 nm   | 352 mm <sup>2</sup> |
| G80             | 681,000,000      | 2006                 | NVIDIA       | 90 nm   | 480 mm <sup>2</sup> |
| R600 Pele       | 700,000,000      | 2007                 | AMD          | 80 nm   | 420 mm <sup>2</sup> |
| G92             | 754,000,000      | 2007                 | NVIDIA       | 65 nm   | 324 mm <sup>2</sup> |
| RV790XT Spartan | 959,000,000      | 2008                 | AMD          | 55 nm   | 282 mm <sup>2</sup> |
| GT200 Tesla     | 1,400,000,000    | 2008                 | NVIDIA       | 65 nm   | 576 mm <sup>2</sup> |
| Cypress RV870   | 2,154,000,000    | 2009                 | AMD          | 40 nm   | 334 mm <sup>2</sup> |
| Cayman RV970    | 2,640,000,000    | 2010                 | AMD          | 40 nm   | 389 mm <sup>2</sup> |
| GF100 Fermi     | 3,200,000,000    | Mar 2010             | NVIDIA       | 40 nm   | 526 mm <sup>2</sup> |
| GF110 Fermi     | 3,000,000,000    | Nov 2010             | NVIDIA       | 40 nm   | 520 mm <sup>2</sup> |
| GK104 Kepler    | 3,540,000,000    | 2012                 | NVIDIA       | 28 nm   | 294 mm <sup>2</sup> |
| Tahiti RV1070   | 4,312,711,873    | 2011                 | AMD          | 28 nm   | 365 mm <sup>2</sup> |
| GK110 Kepler    | 7,080,000,000    | 2012                 | NVIDIA       | 28 nm   | 561 mm <sup>2</sup> |
| RV1090 Hawaii   | 6,300,000,000    | 2013                 | AMD          | 28 nm   | 438 mm <sup>2</sup> |
| GM204 Maxwell   | 5,200,000,000    | 2014                 | NVIDIA       | 28 nm   | 398 mm <sup>2</sup> |
| GM200 Maxwell   | 8,100,000,000    | 2015                 | NVIDIA       | 28 nm   | 601 mm <sup>2</sup> |
| Fiji            | 8,900,000,000    | 2015                 | AMD          | 28 nm   | 596 mm <sup>2</sup> |

2007

# Paul S. Otellini Intel Corporation's fifth CEO



# Why are we continuing to strive for smaller and smaller technology?

---

- More transistors/chip → increased functionality and performance
- Higher speeds → partially depends on how close together the components are placed
- Cheaper – more chips/wafer, greater yields

# Yield Ratio

---

$$yield = \frac{n_w}{n_t}$$

$$n_w = yield \bullet n_t$$

$n_w$  = number of working chips/wafer

$n_t$  = total number of chips/wafer

Old fab lines, yield  $\rightarrow > 90\%$

New fab lines, yield  $\rightarrow < 40\%$

## A circular grid of 10x10 squares, with a diamond-shaped pattern of red and gray squares. The grid is centered on a white circle, which is set against a dark brown background. The pattern consists of red squares forming a diamond shape, with gray squares filling the rest of the grid. The red squares are located at the following coordinates (row, column) starting from the top-left corner: (1,4), (2,3), (2,5), (3,2), (3,6), (4,1), (4,7), (5,1), (5,2), (5,3), (5,4), (5,5), (5,6), (5,7), (5,8), (5,9), (6,2), (6,3), (6,4), (6,5), (6,6), (6,7), (6,8), (6,9), (7,3), (7,4), (7,5), (7,6), (7,7), (7,8), (7,9), (8,4), (8,5), (8,6), (8,7), (8,8), (8,9), (9,5), (9,6), (9,7), (9,8), (9,9), (10,6), (10,7), (10,8), (10,9). The gray squares are located at the following coordinates: (1,5), (2,4), (2,6), (3,3), (3,7), (4,2), (4,8), (5,2), (5,3), (5,4), (5,5), (5,6), (5,7), (5,8), (5,9), (6,1), (6,2), (6,3), (6,4), (6,5), (6,6), (6,7), (6,8), (6,9), (7,1), (7,2), (7,3), (7,4), (7,5), (7,6), (7,7), (7,8), (7,9), (8,1), (8,2), (8,3), (8,4), (8,5), (8,6), (8,7), (8,8), (8,9), (9,1), (9,2), (9,3), (9,4), (9,5), (9,6), (9,7), (9,8), (9,9), (10,1), (10,2), (10,3), (10,4), (10,5), (10,6), (10,7), (10,8), (10,9).

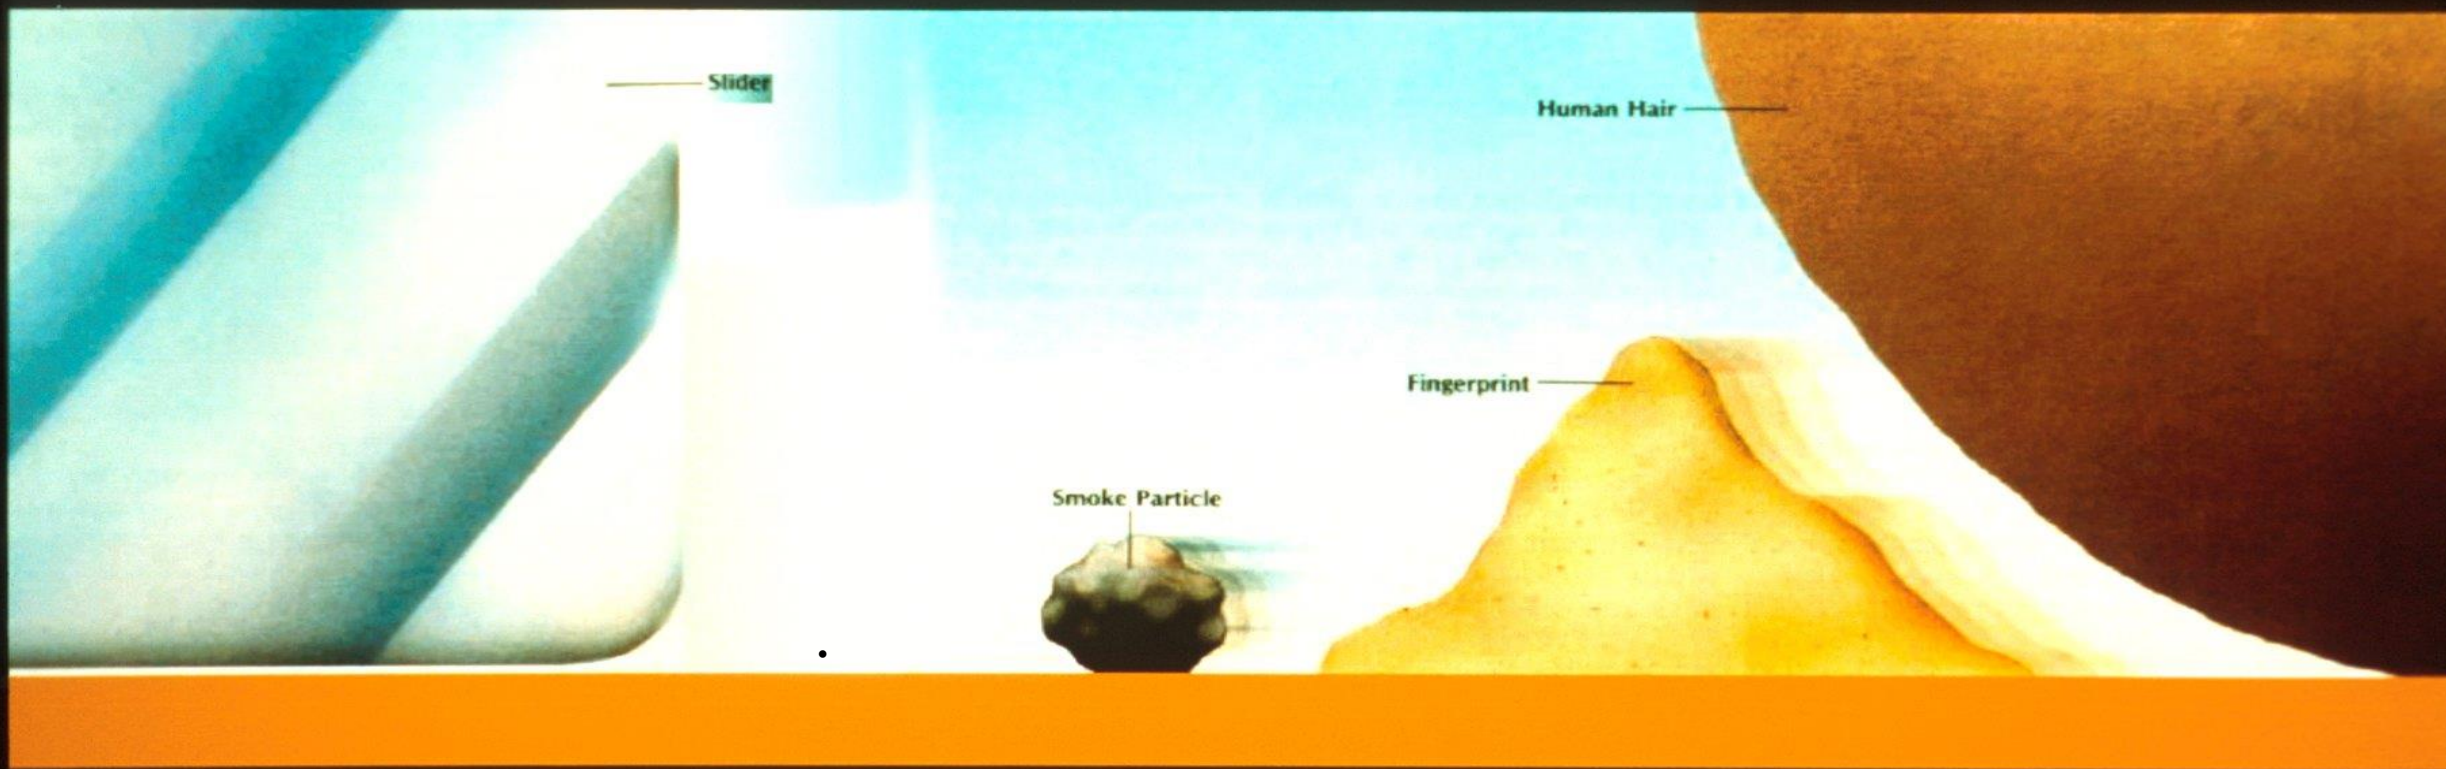
# Yield Ratio

---

Number of defects/unit area depends on the process

$$\therefore \text{Yield} \approx \frac{1}{\text{Chip Area}}$$

Total chips ( $n_t$ ) for a given wafer size is also inversely proportional to the chip area



---

**Why does the shrinking technology make  
the cost of manufacturing cheaper per  
component?**

# Example:

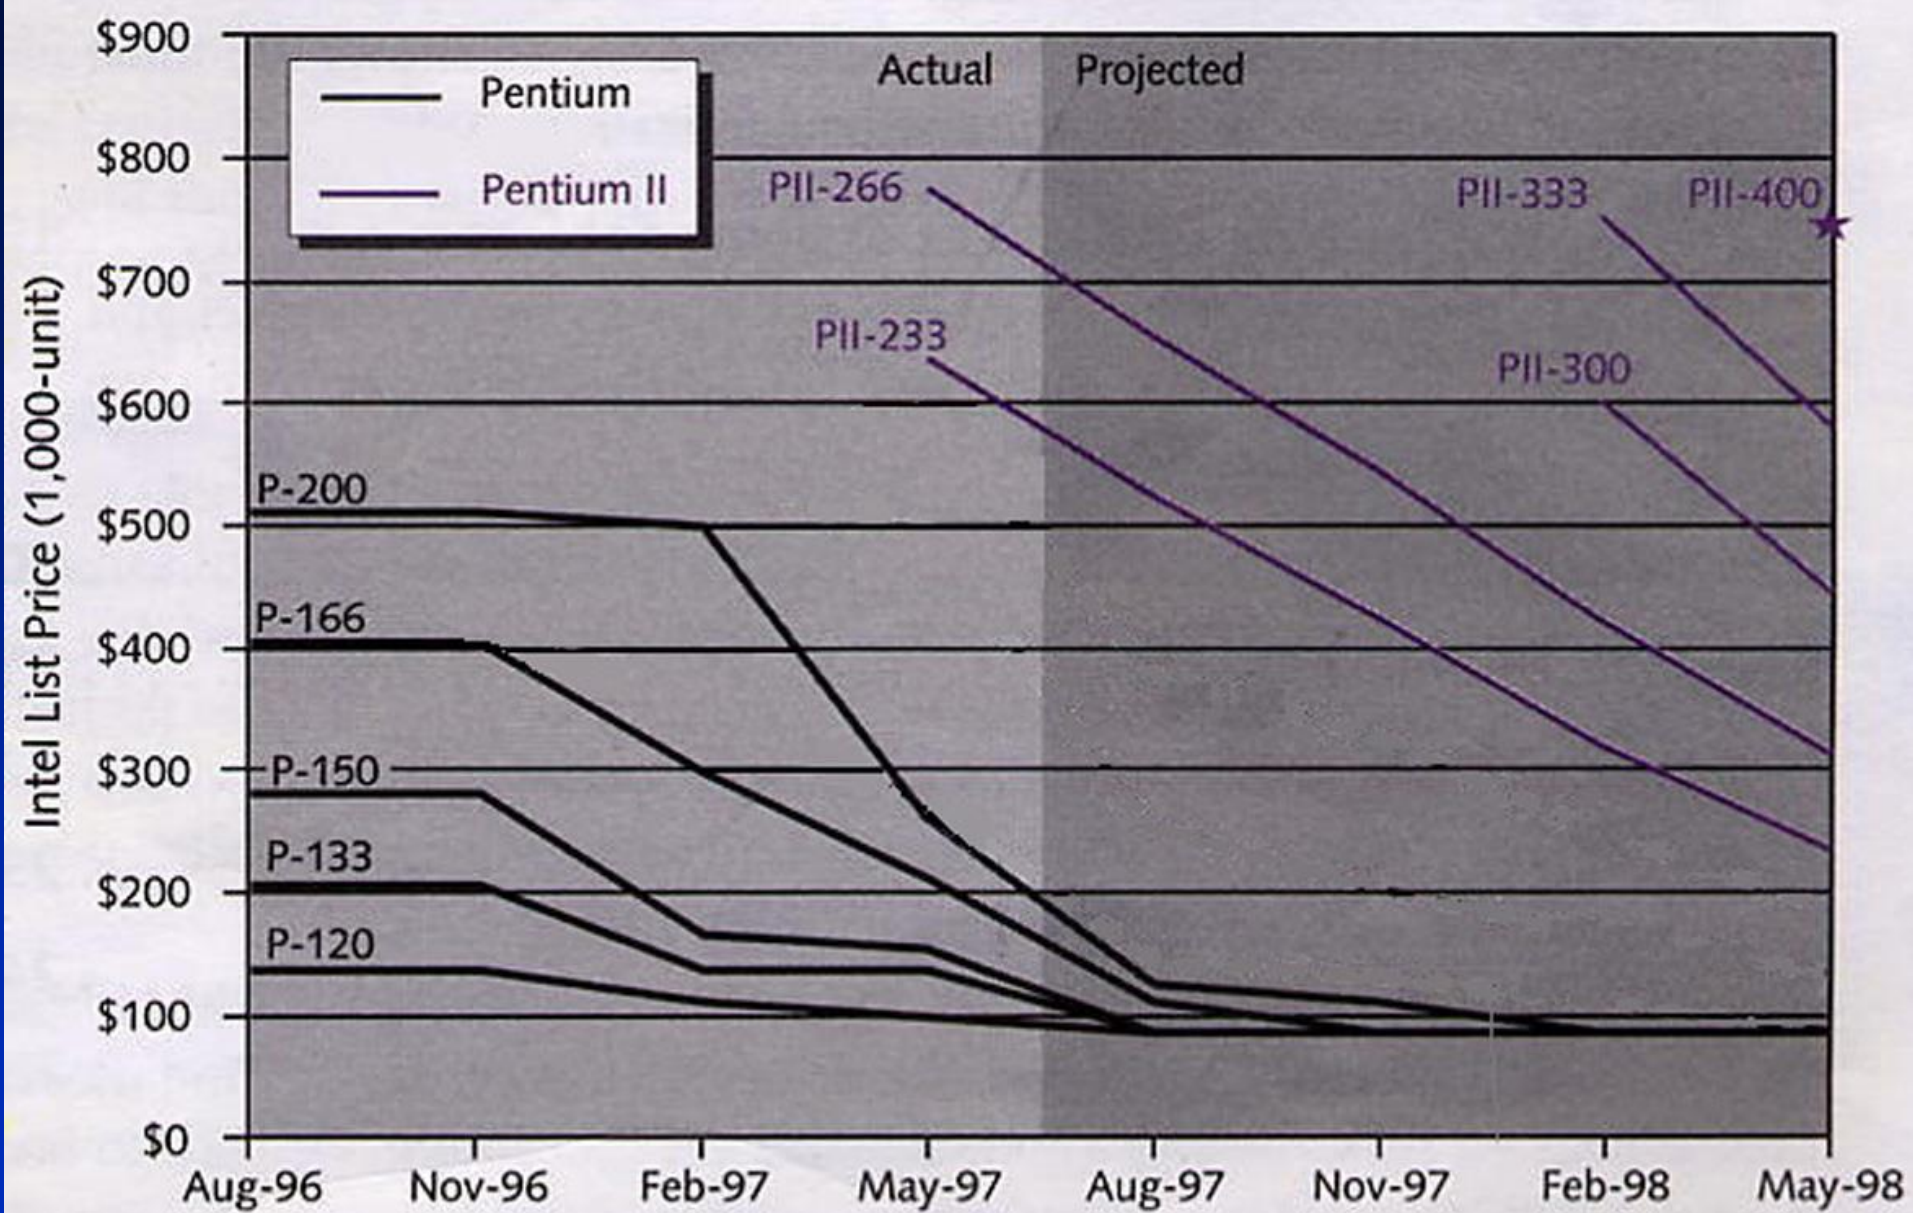
For a 10% shrink in feature size :

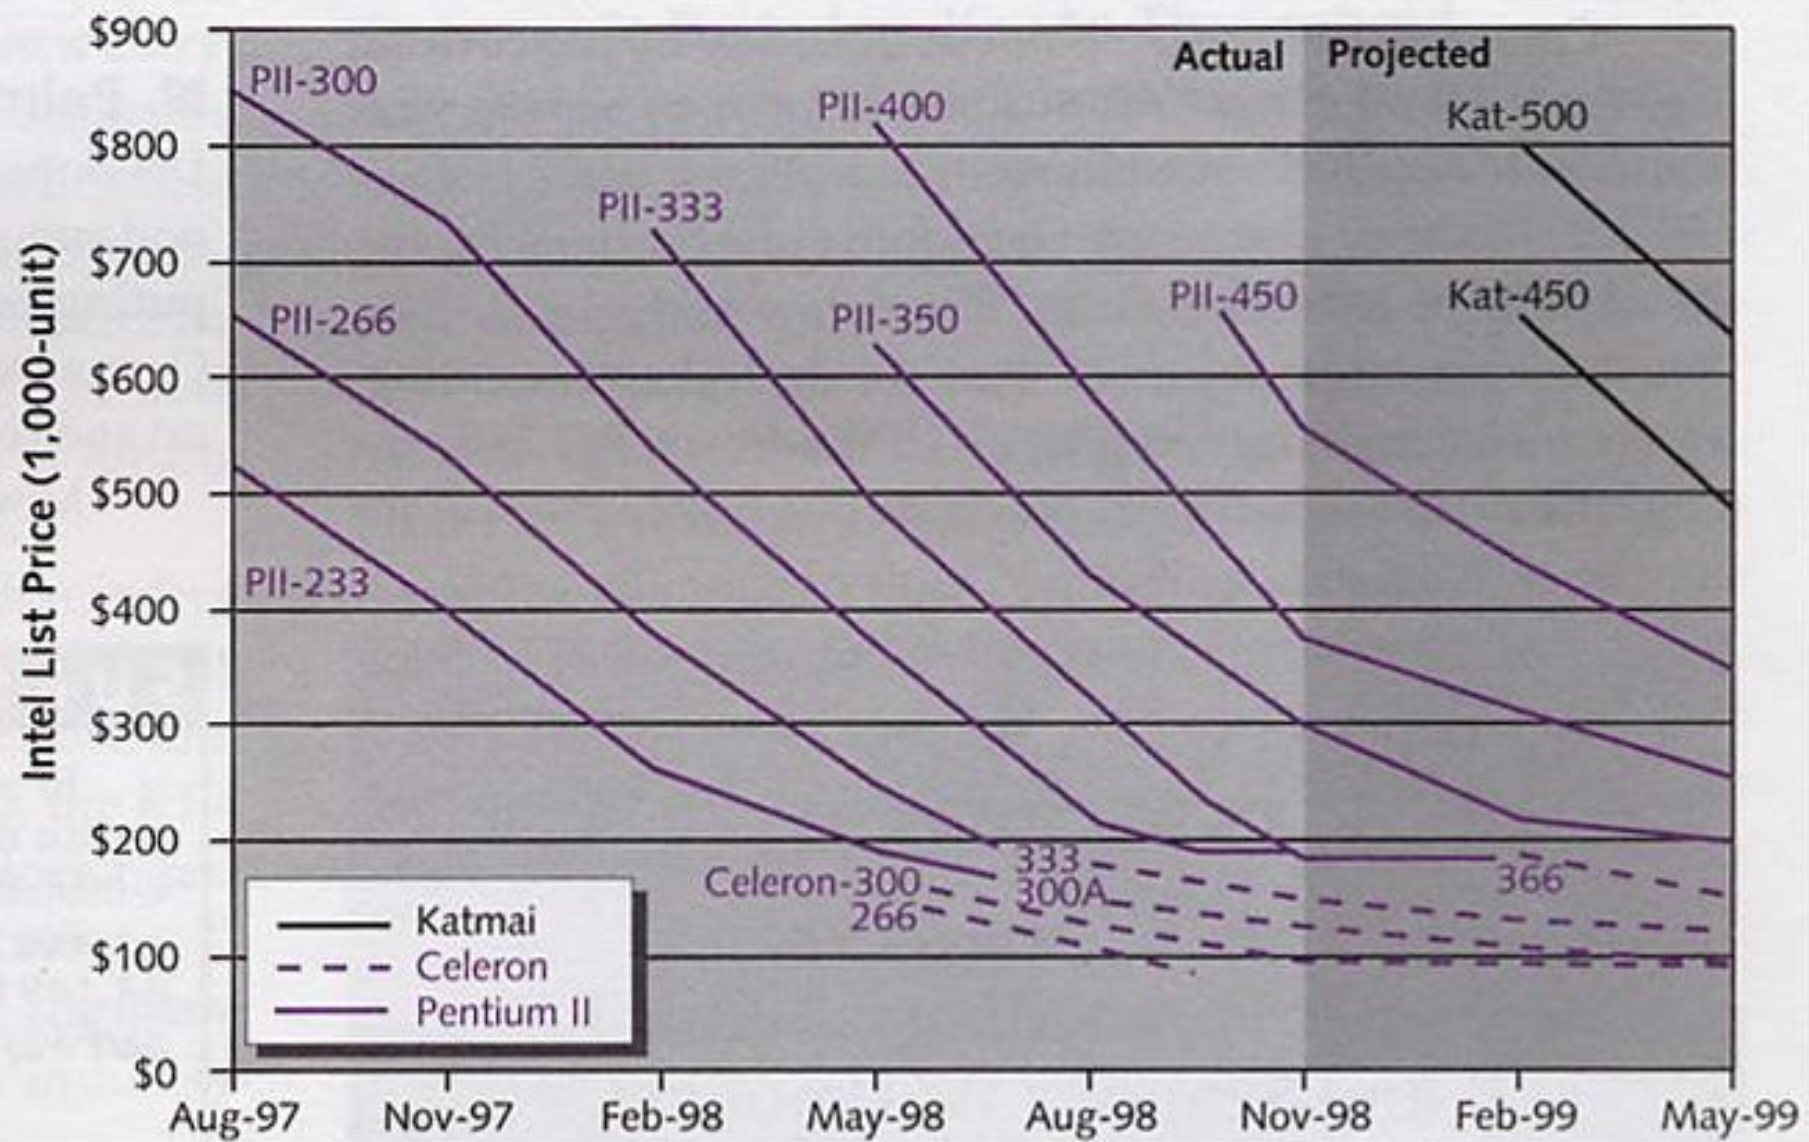
$$n_{w_{new}} = n_{w_{old}} \left( \frac{1}{.9} \right)^2 \left( \frac{1}{.9} \right)^2$$

↑
↑

New yield
New  $n_t$

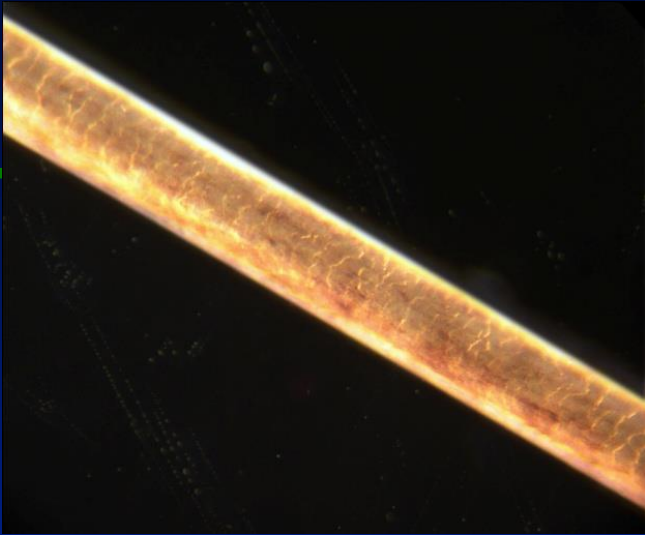
$$n_{w_{new}} = 1.52n_{w_{old}}$$





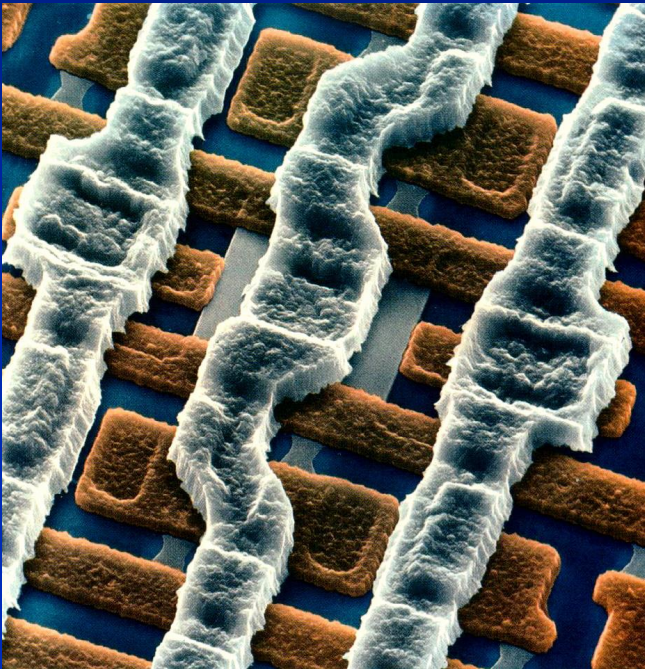
---

**Can this  
Shrinking Technology  
Continue?**



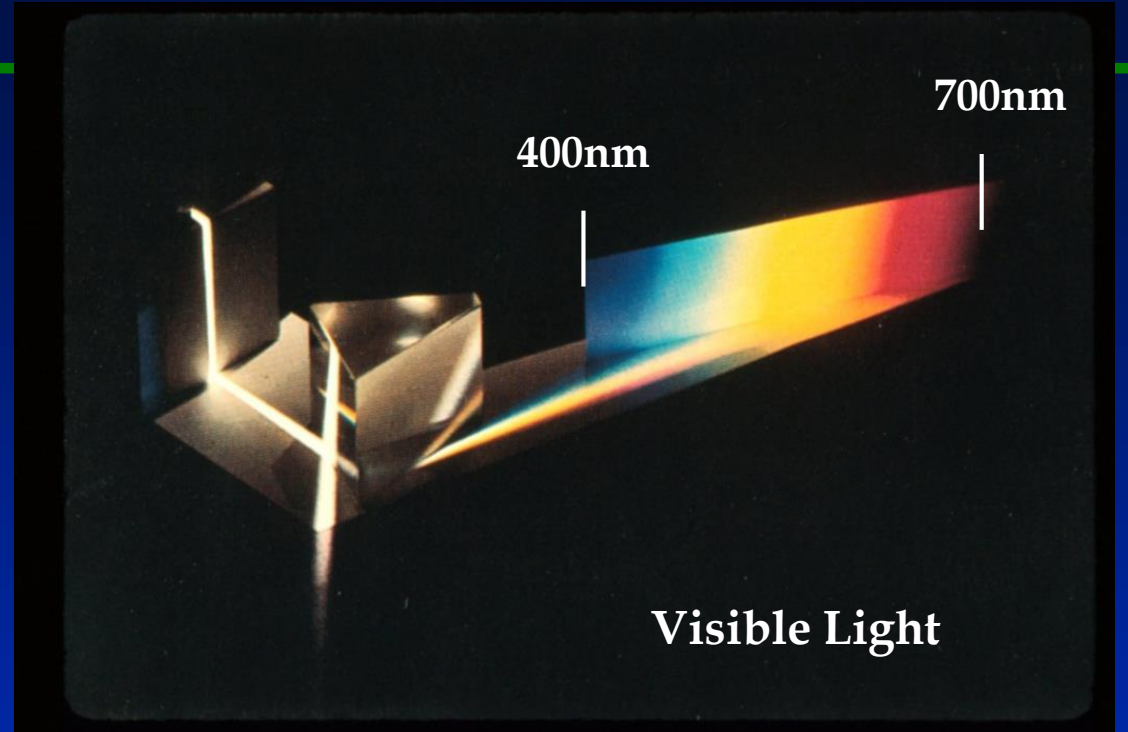
Human Hair

50 $\mu$ m  
(50,000 nm)

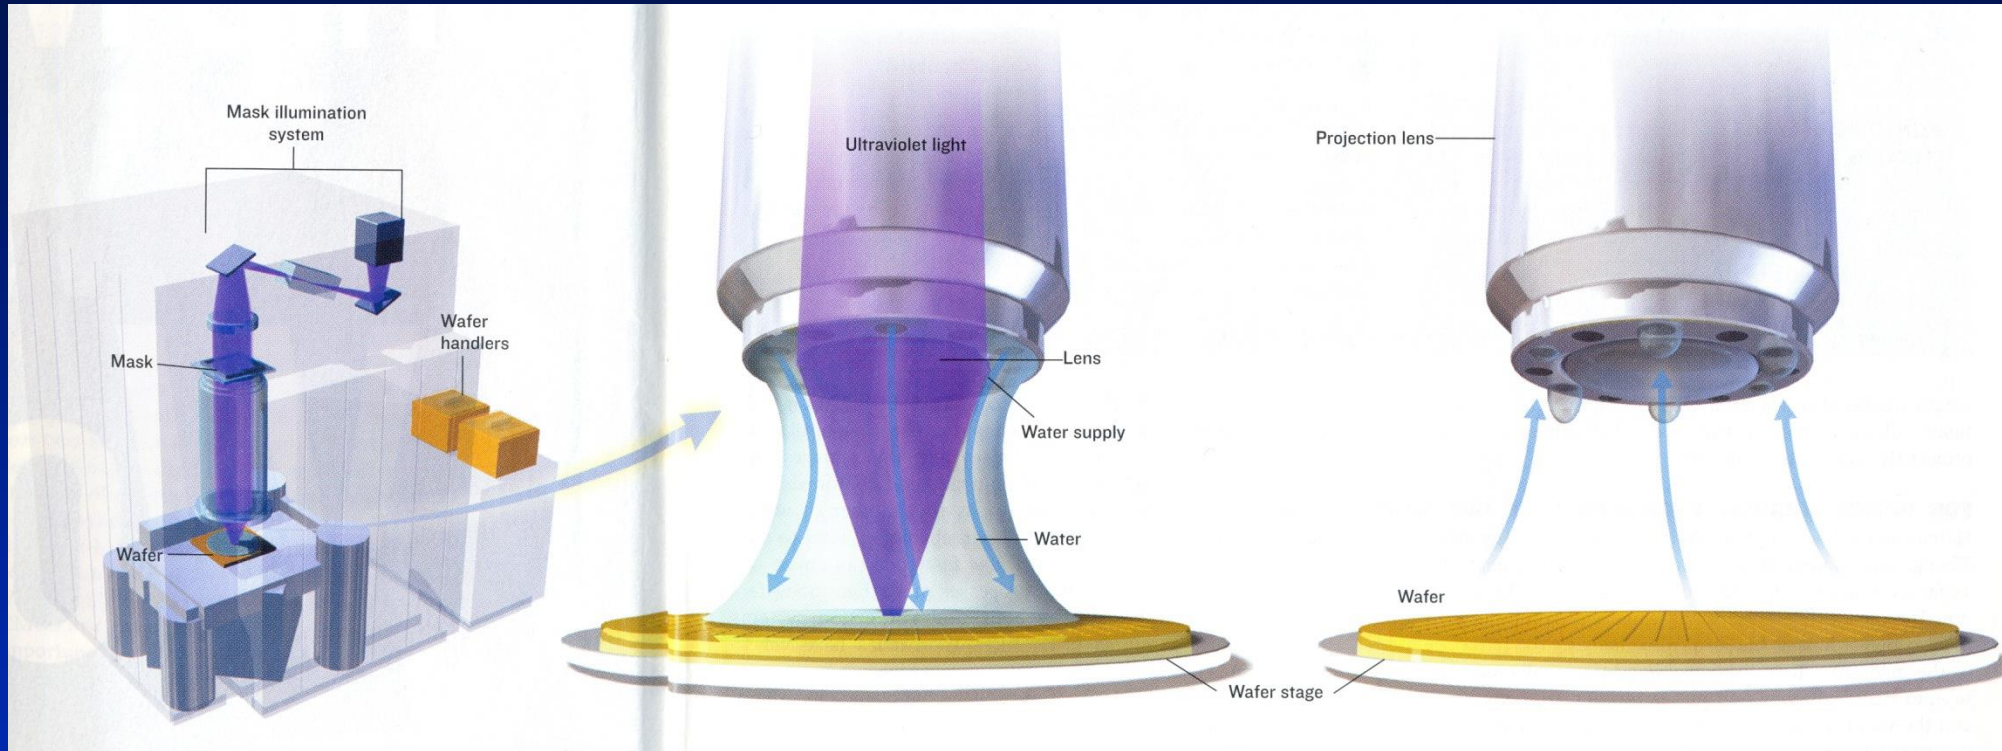


Integrated Chip

32nm



# Getting Wafers Wet



By adding a thin layer of water between the projection lens and the wafer, the immersion system can create features 30 percent smaller.

# Photolithography

---

- Defining the smallest components requires short wavelengths of light.
- Currently, most fabrication processors use extreme ultra-violet light at 193nm.
- Can pass the light through water. The water slows the light (less velocity) shrinking its wavelength. It is estimated that this technique will meet demands for 7 more years.
- On February 20, 2006 IBM Almaden & JSR Micro demonstrated a system using an “unidentified” light slowing liquid yielding patterns 29.9nm wide.

# The HIGH-k SOLUTION

By Mark T. Bohr, Robert S. Chau, Tahir Ghani  
& Kaizad Mistry — October 2007 IEEE Spectrum

In 2007 new 45nm  
Microprocessors were the  
result of the first big material  
redesign in CMOS transistors  
since the late 1960s





# 10 Nanometer Technology

---

- Nov. 15, 2012, Samsung unveiled a 64 gigabyte (GB) multimedia card (eMMC) based on 10 nm technology.
- April 11, 2013, Samsung announced it was mass-producing High-Performance 128-gigabit NAND Flash Memory with 10 nm and 20 nm technology.
- April 2015, TSMC announced that 10 nm production would begin at the end of 2016.
- May 23rd 2015, Samsung Electronics showed off a wafer of 10nm FinFET chips.

# Factors Contributing to Advancing Microprocessor Performance

---

- Shrinking Component Size
- Increasing Speed
- Reducing Circuit Resistance
- New Materials

# Factors Contributing to Advancing Microprocessor Performance

---

- RISC vs. CISC
- VLIW
- Multi-level Cache
- Parallelism & Pipelining

# Factors Contributing to Advancing Microprocessor Performance

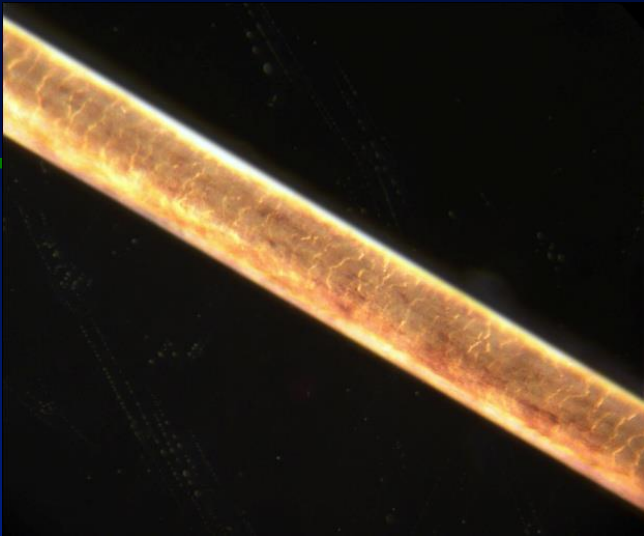
---

- RISC vs. CISC
- VLIW
- Multi-level Cache
- Parallelism & Pipelining
- Multi-core Technology

# Multicore Craze

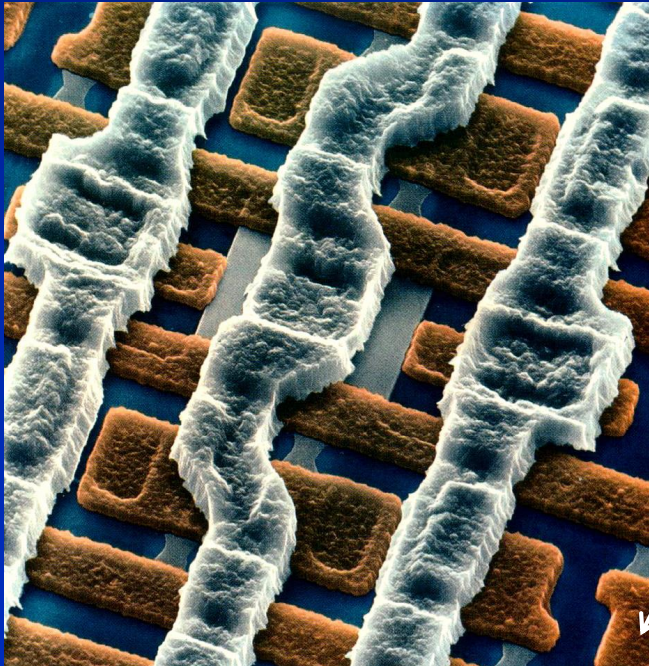
---

- For years, the trend was to make chips faster  
Today → 3 Ghz
- But the power required (Watts) and the heat generated is proportional to the frequency squared.
- Therefore, put more computers on the chip but run at slower speeds.



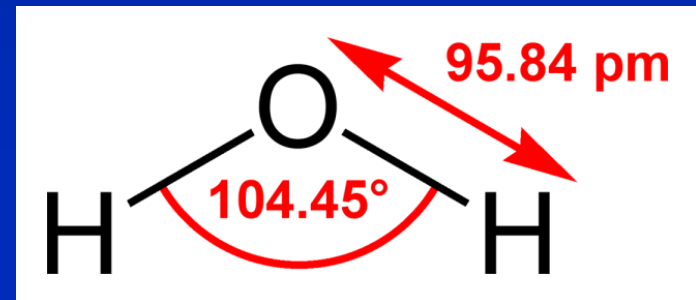
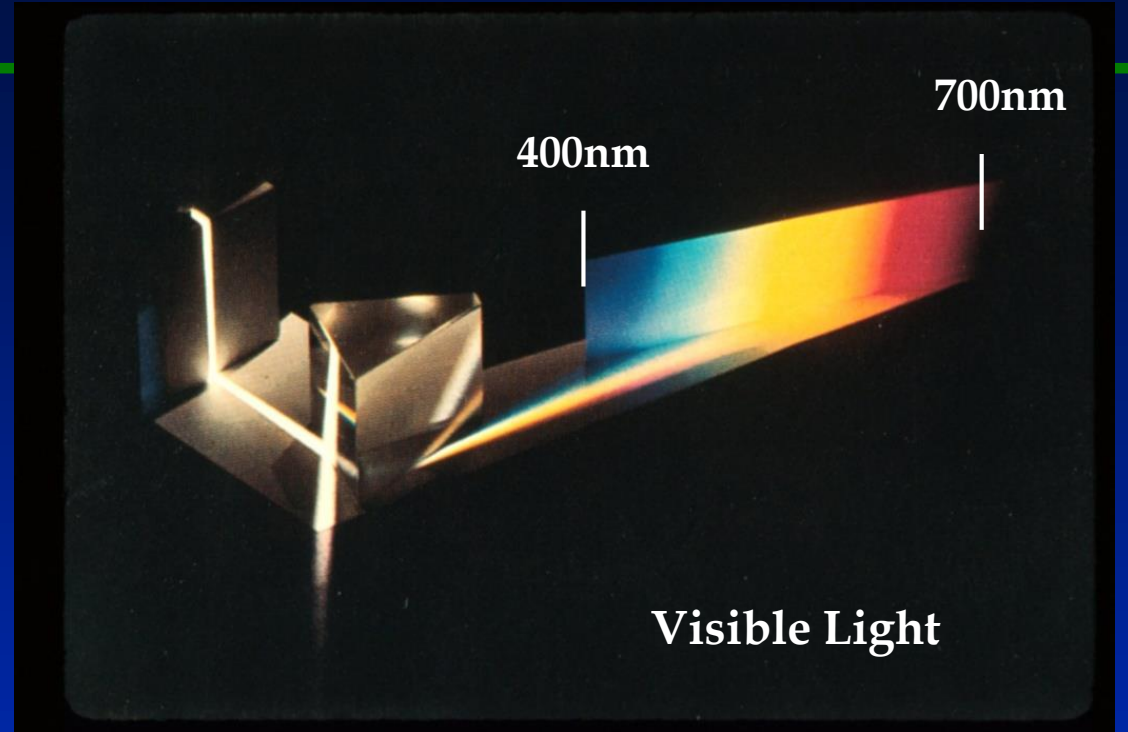
Human Hair

50 $\mu$ m



Integrated Chip

32nm



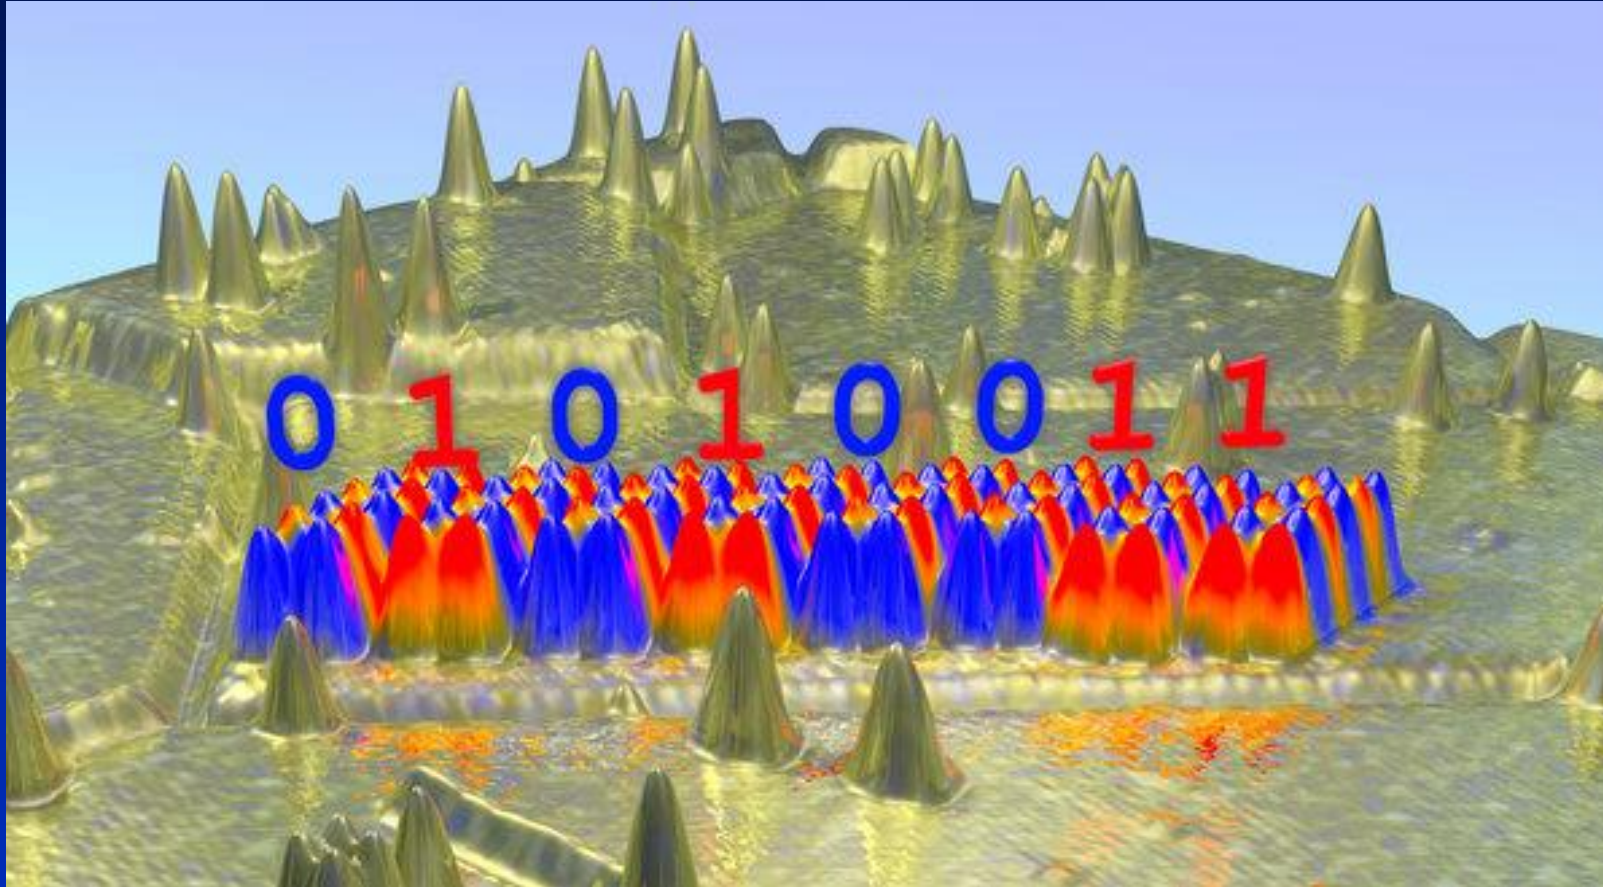
- 
- How long can Moore's Law continue?
  - What are the limits to this integrated circuit technology?

*“There are two constraints:*

- The finite velocity of light*
- The atomic nature of materials”*

- Stephen Hawking

# Miniaturized Data Storage at Atomic Scale

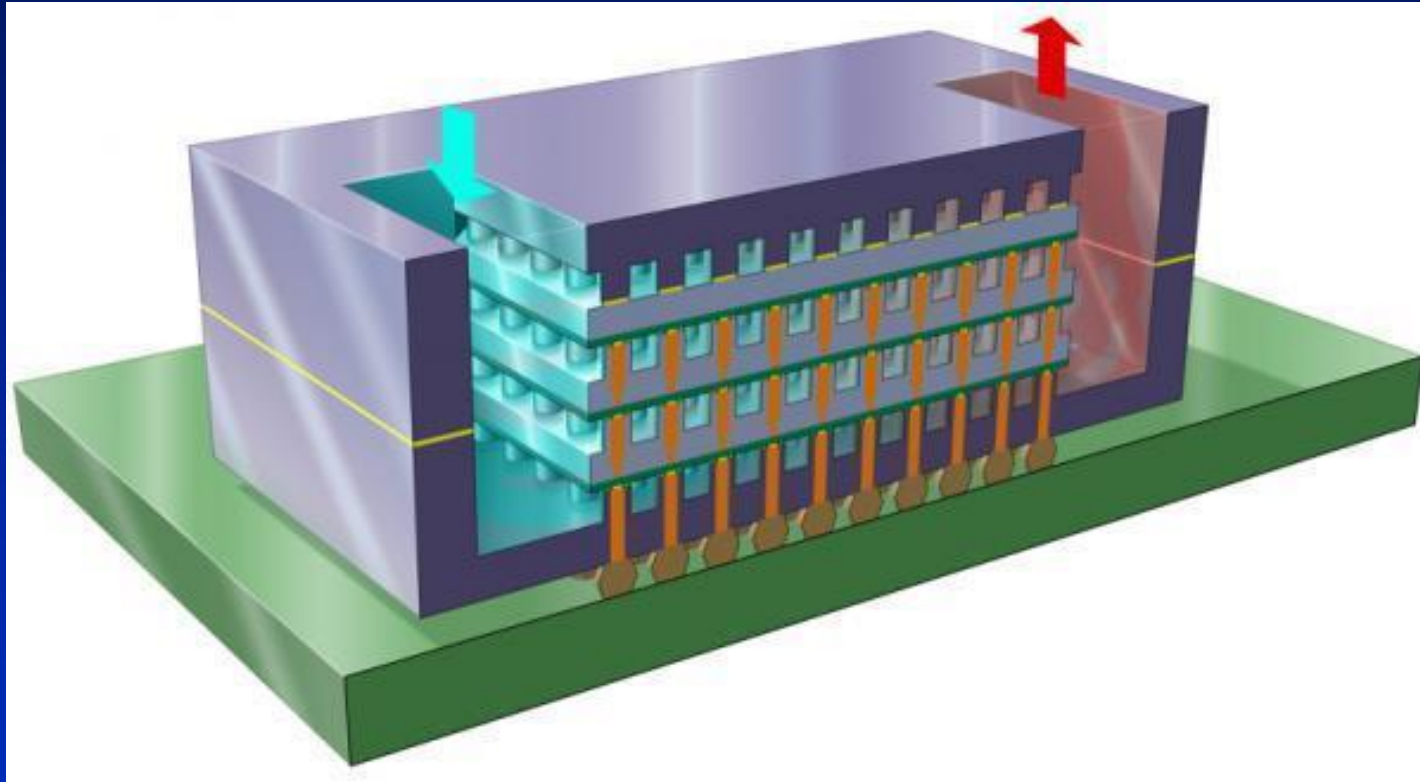


IBM researchers have stored and retrieved digital data from an array of just 12 atoms

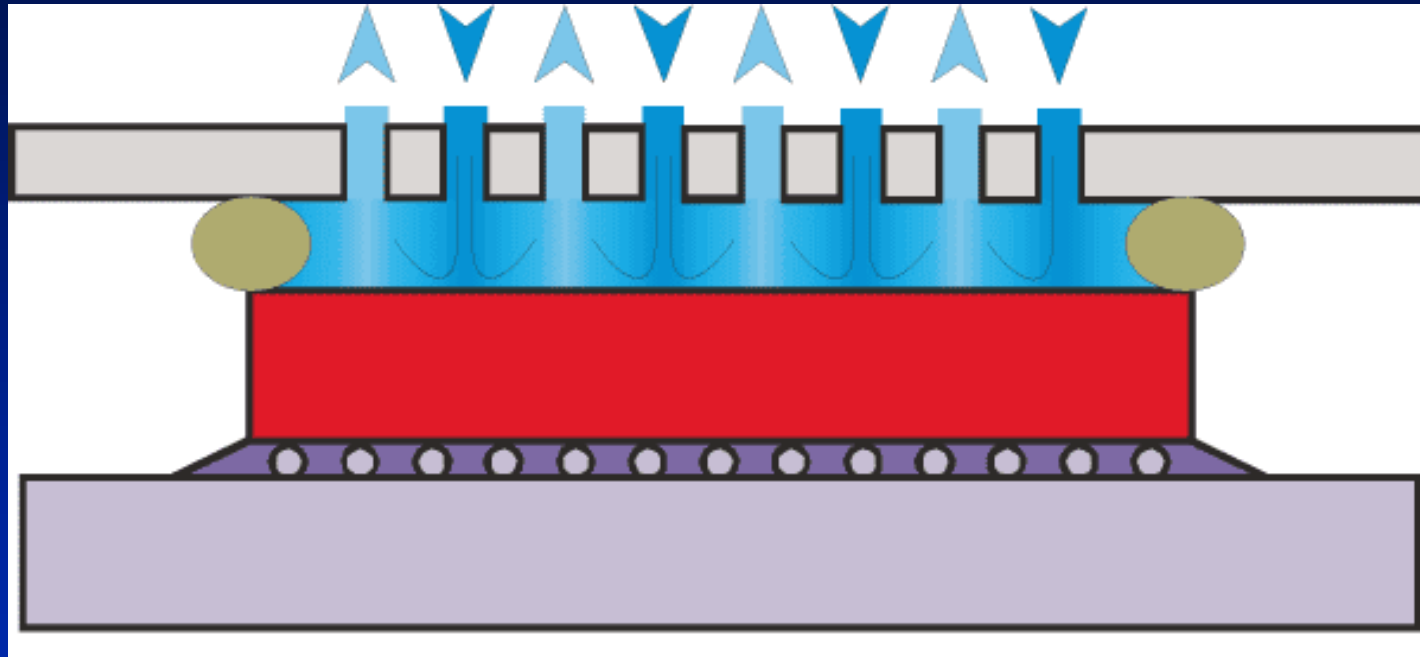
New York Times, 1/12/12

# IBM's Chip Stacking Technology

---

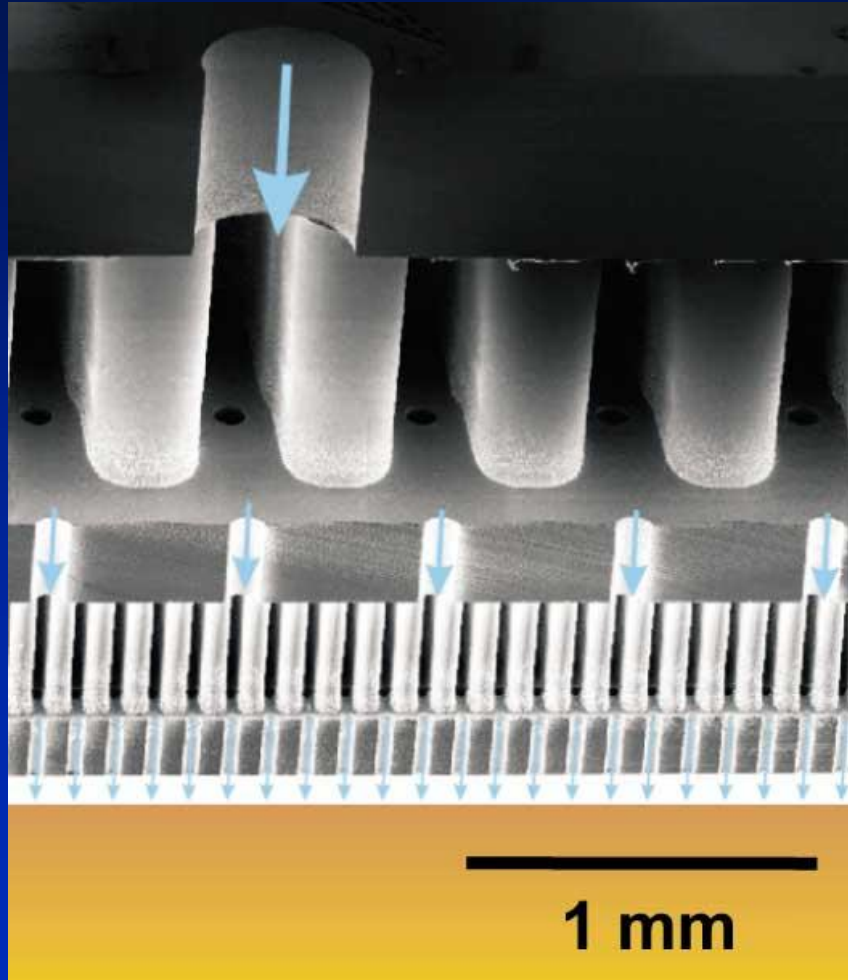


# Single-phase, miniaturized convective cooling



Distributed return architecture with cross section showing inlet jets with neighboring drainage holes.

# Single-phase, miniaturized convective cooling



SEM section of two-level jet plate. Water flow is indicated by blue arrows.

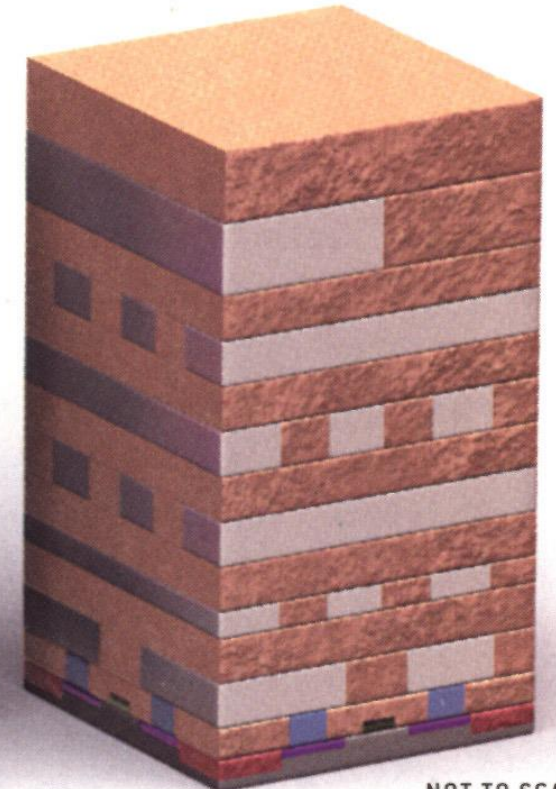
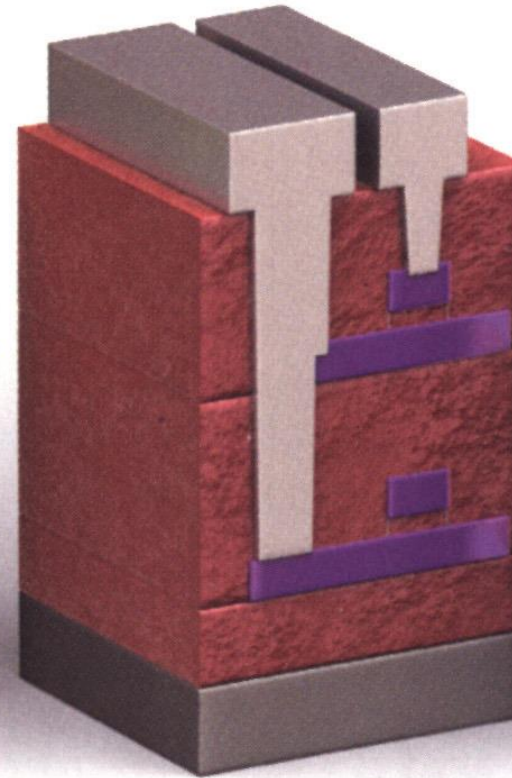
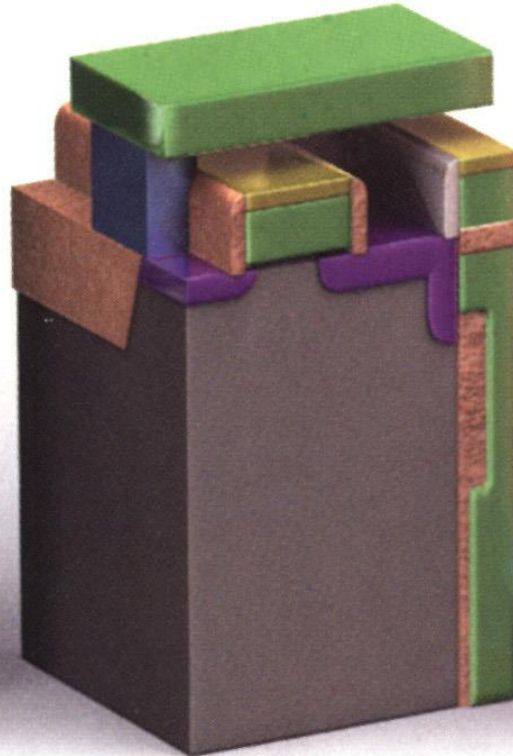
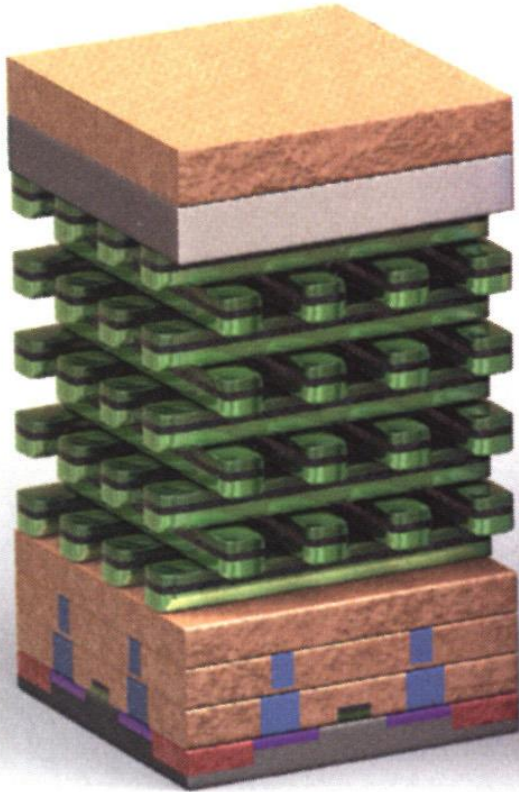
# Interior Structure of 3-D chips

3-D Volatile Memory  
[Matrix Semiconductor]

2-D Random-Access Memory  
[IBM 256-Megabit]

3-D Logic Circuit  
[Lab Prototype]

2-D Microprocessor  
[Advanced Micro Devices Athlon]



NOT TO SCALE

■ Monosilicon substrate   ■ Insulators   ■ Aluminum wires   ■ Polysilicon   ■ Tungsten plugs   ■ Ion-doped silicon   ■ Isolation oxides   ■ Silicide

# Intel's 3D Transistor

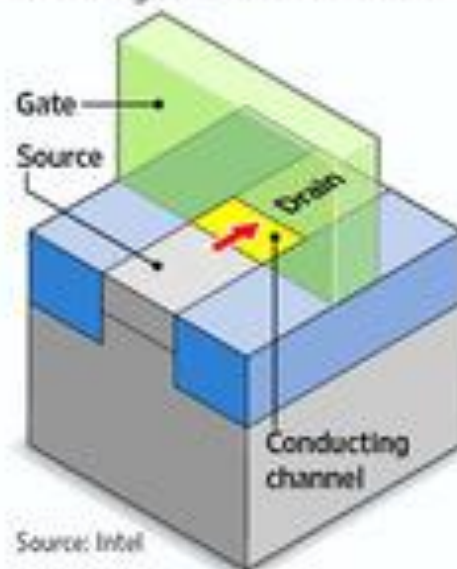
2011

## Intel's Move Into 3-D

The chip maker breaks from conventional approaches to make transistors.

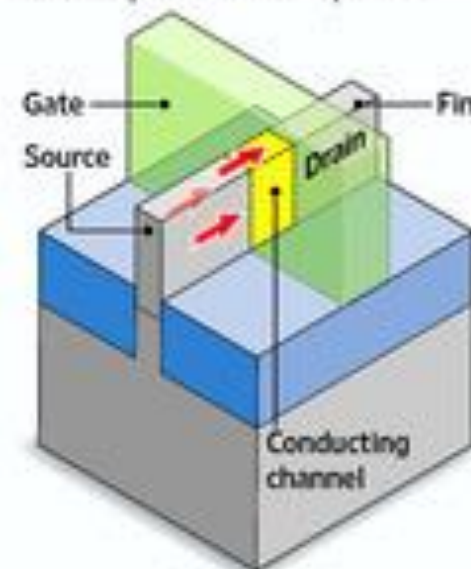
### Conventional transistor:

Electrons flow between components called a **source** and a **drain**, forming a two-dimensional **conducting channel**. A component called a **gate** starts and stops the flow, switching a transistor on or off.



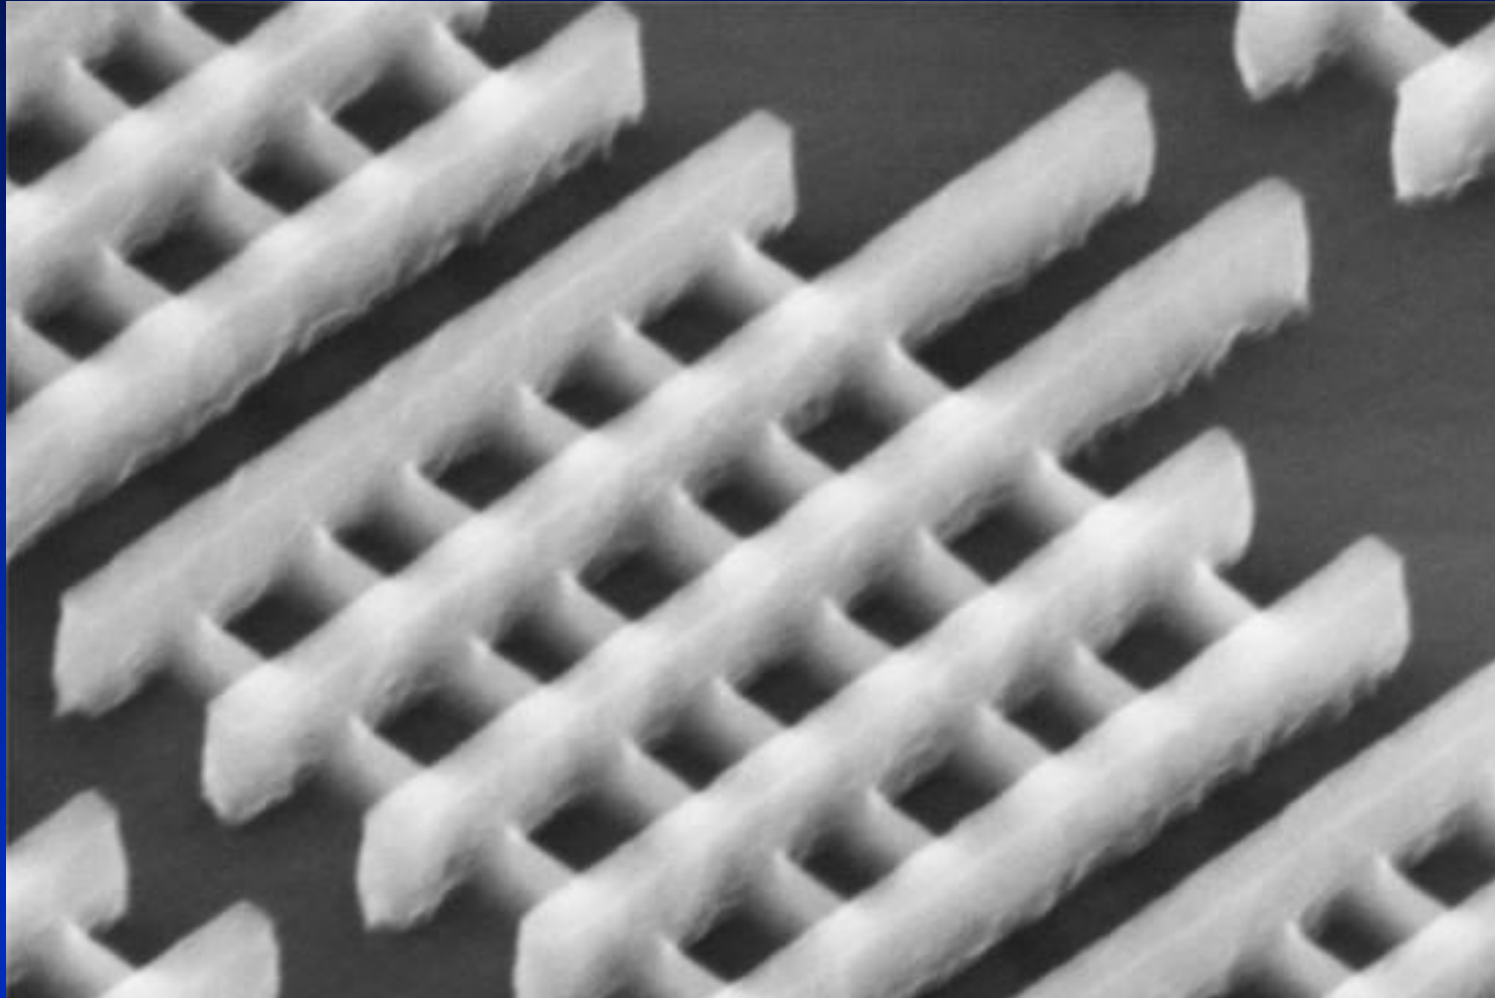
Source: Intel

**Intel's new transistor:** A fin-like **structure** rises above the surface of the transistor with the **gate** wrapped around it, forming **conducting channels** on three sides. The design takes less space on a chip, and improves speed and reduces power consumption.

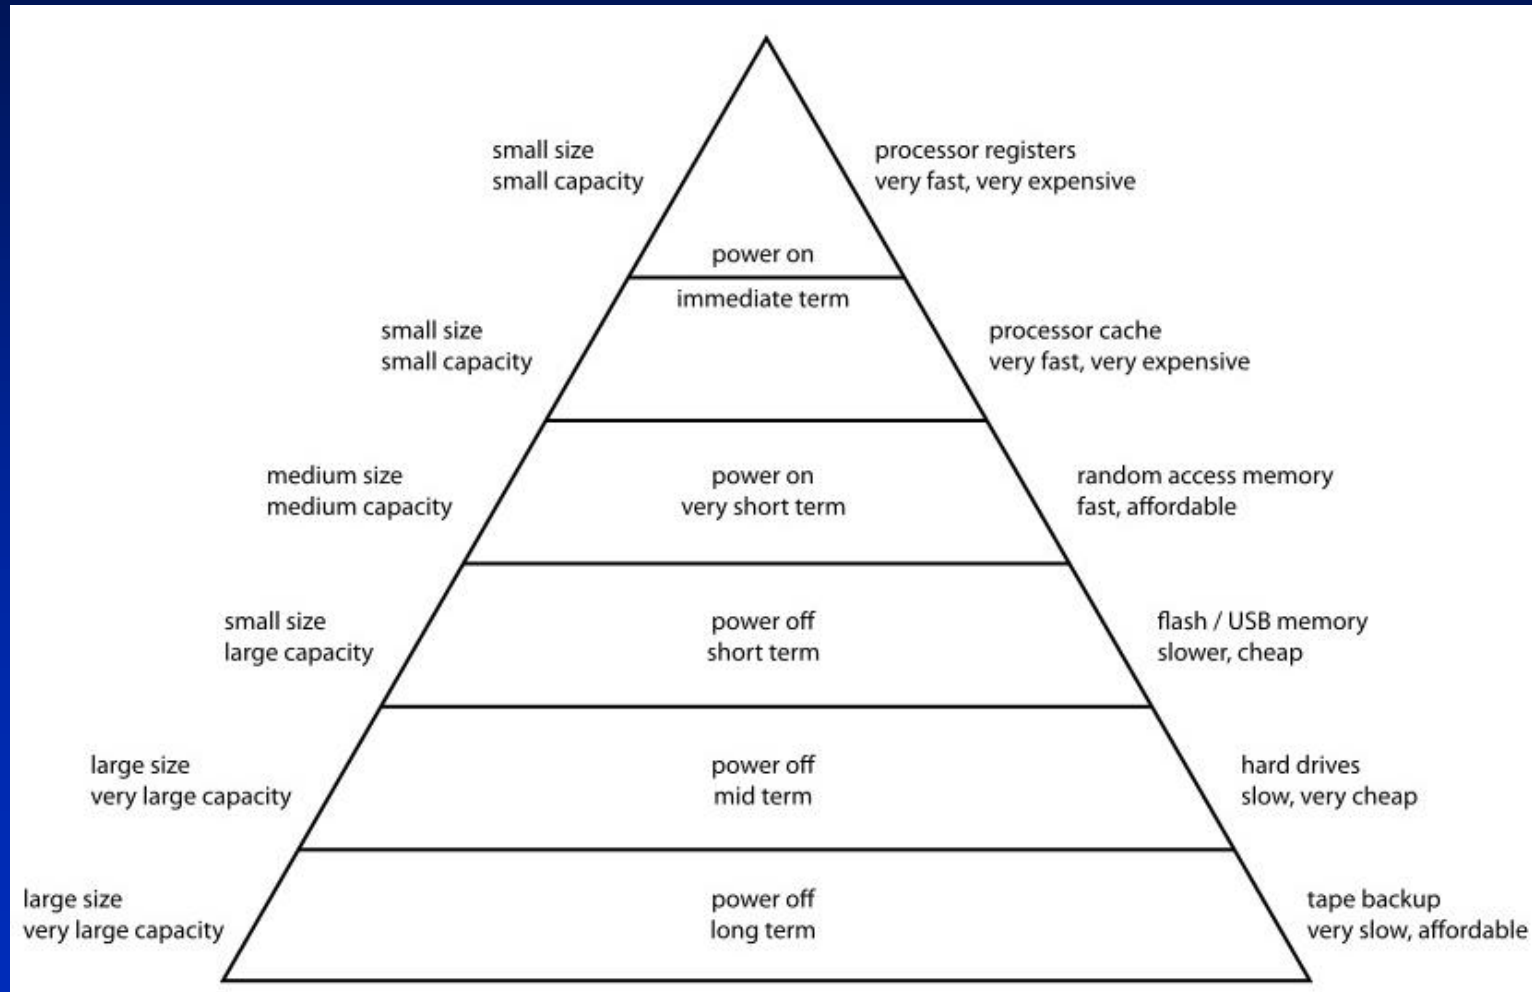


# Intel's 22nm 3D tri-gate transistor

---



# Computer Memory Hierarchy



---

“Every economic era is based on a key  
abundance and a key scarcity.”

*George Gilder,  
Forbes ASAP, 1992*

What are the key scarcities?

# End

---